

Probing the higher redshift universe by studying strong
lensing of gravitational waves and enhancing search
sensitivity of the GstLAL search pipeline

Thesis by
Alvin Ka Yue Li

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy in Physics



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2024
Defended May 3, 2024

© 2024

Alvin Ka Yue Li

ORCID: [0000-0001-6728-6523]

All rights reserved

ACKNOWLEDGEMENTS

To begin, I would like to dedicate my thesis to

- **my most respected teacher, Professor Alan J. Weinstein,**
- **my dearest brother, Elwin Li,**
- **my loving mother, Wendy Choy,**
- and, to answer Alan's question on the original title of my thesis, **the anime/-manga GIVEN (ギヴン) written by Natsuki Kizu (キズナツキ).**

I believe many people will think that I am eccentric (it is already a much better word than crazy) to dedicate my thesis to a Japanese anime. Please bear with me and allow me to explain. Particularly, “GIVEN” is an inseparable part of my PhD journey. I resonate with “GIVEN” firmly.

Let me try to present it this way: In the following, boxed texts are summaries of “GIVEN”, and unboxed sections are stories of mine. Alan knows part of the story already, and maybe my friend and colleague Simona has also heard me talking partially about it back when I was a SURF, but I have yet to tell them the full version of the story. This is the right time.

┌ The story of “GIVEN” revolves around the main character, Mafuyu Sato (佐藤真冬). Mafuyu was born into a typical family, living with his father and mother. However, he was constantly abused (verbally and physically). Mafuyu explained in one of the scenes that “his father would beat him if he spoke” (喋ると、お父さんが打つの). All these cause Mafuyu to become a depressed and quiet boy. Only when he met another boy one day, Yuki Yoshida (吉田由紀), who listened to him and called the police for help, did Mafuyu manage to escape from his father's violence and regain happiness. Ever since Mafuyu's father was arrested, Yuki promised to be with and protect Mafuyu forever (俺が一緒にいる、ずっとだ). └

I was born into an ordinary family in 1997, living with my father and mother. I would have had a great childhood, I really would have, but it turns out I did not. As I remember, I used to live in a decent apartment, which my mother and father were

in the process of purchasing through installments. However, one day, we moved out of the apartment to a rental apartment, and then we moved again, one after another, to smaller and smaller apartments. Back then, I had no idea what was happening. When I turned five, my brother Elwin was born and entered the family. Curiously, however, my father spent less and less time at home. My mother used to take Elwin and me to meet many of her friends and our relatives, but that frequency also decreased eventually. Again, back then, I had no idea what was happening.

2004 was the first turning point of my life that I will never forget. I was a Primary 2 student (equivalent to Grade 2 in the US). I remember I got “first in form” in the final exam at the end of the school term. It should have been a happy news to celebrate. However, it was at the very same time that my mother discovered that my father was having an affair. They had a big fight, and a divorce was inevitable. The whole thing was devastating to the seven-year-old me. But that was not the worst part.

My father (hereafter, I will refer to him as that man because he does not deserve to be my father) left me and my family with one final statement: “Listen, Alvin, me having the divorce with your mother is the best present I will give you for your birthday ¹ and for you getting first in form in your final exam.” It was a very cruel statement indeed, especially to a young child like me. I must have been crying hard at that time, and I believe I even felt that all my effort was worthless. But my mother, who should have been much more depressed than I was, hugged me tightly and cheered me up. Regardless of what she was going through back then, she toughened herself and continued to take care of and protect Elwin and me by all means as both a loving mother and a strict father. She promised to stay with us forever. Since then, my goal has been to work hard to make a good showing for my mother.

Eventually, Mafuyu and Yuki became boyfriends. Yuki lived up to his promise - he gave Mafuyu his unlimited love and care. They were together all the time, chatting, playing, spending time together. Mafuyu relied on Yuki, and Yuki relied on Mafuyu. It would have been a great story for the two of them.

My mother sacrificed everything to take care of me and Elwin. Back then, that man not only stole my mother’s money (which is why we had to go through the constant apartment downsizing). He even borrowed money from our relatives and

¹ That man left us in July 2004; my birthday is August 17.

my mother's friends, using false excuses that we did not have enough money for food and lying that my mother was unwilling to buy milk powder for Elwin, who was still a baby then. Yes, such a terrible man does exist on Earth, and from my point of view, he is not even worthwhile to be called a human. At that time, my mother must have been ashamed, and she must have been embarrassed to face her friends and our relatives. She must have suffered a lot. She must have been deeply depressed. However, even in such a horrible situation, she never considered giving us up.

My mother gave up her dignity to apply for "Comprehensive Social Security Assistance" (financial assistance) from the government. She could have applied for a decent job, given that she was an assistant manager previously before resigning to become a full-time housewife. But she decided it was more important to spend unlimited time caring for me and Elwin, especially because we were still recovering after the incident.

My mother stayed strong and took up both the role of a mother and a father to nurture us. She was very strict with us, and that is because she was, at the same time, a father to us, and she had to make sure we stayed on the right path. Yes, sometimes we felt that she was too strict, and sometimes she would yell at us for no reason, and we did complain about it. But as I grew up, I gradually understood why she had to be that strict, and I also understood the pressure my mother was bearing, and that we should be more considerate. I will never forget one thing my mother always told me: "Remember Alvin, I would not have scolded you that hard if I did not care about you. It takes much energy to scold and be angry with someone, and it makes one feel sad."

My mother never cared about herself. You would notice that she always wore the same clothes and shoes and used the same handbag over the years. If you looked closer, you would have noticed the holes within and the many signs of repairing on them. That is because she was never willing to spend money on herself. She was willing to spend money on us to buy new clothes and shoes, but she would only allow herself to buy HK\$20 shoes from the bargain shops in Mong Kok. She did not even care if the shoes did not fit her foot size or were hurting her feet. For the shoes at home, even if they were broken, she would only use glue to repair them, and only when they were no longer repairable would she throw them away. She would rather use super glue to fix her broken glasses than to buy a new pair; when Elwin and I bought breakfast from fast-food shops, she would never join us in purchasing

a meal; instead, she would just eat the toast from our meals. She never stopped us from trying all the street food we desired when we went out. Even though our family is poor, she encouraged Elwin and me to join all sorts of study tours and summer research exchange programs. Who knows how much she had given up to fulfill our many needs and desires?

My mother had always protected us from all kinds of attacks, verbal or physical. I still remember the mocking from my relatives like “what is the big deal about getting into university?”, “You should be grateful that an institute in the US is willing to take your son as a graduate student”... My mother always stood before us as a giant shield to protect us from these mockeries. And that is why I promised myself to work hard to win credit for her, hoping that one day she will be proud of us, that she can take a good rest and enjoy life.

Nine years ago, when I became an undergraduate, and Elwin was still a high school student, the heartless government decided to cut down the subsidy to our family to force my nearly 50-year-old mother to return to the workforce. My mother did not work in the first place to get the most time to take care of us. That was terrible timing because Elwin was also going through HKDSE (similar to the university entrance exam), one of the most critical exams in his life, and it was when he needed my mother’s support the most. My mother decided to take on a job that she did not like - being a security guard, only because she could leave work early and get home early so she could still take care of Elwin. Being a security guard definitely differs from a job she liked, but she worked hard and gradually gained her boss’s trust. Patrolling day after day must have been a significant burden to her, but she never ditched her role as a mother, even though she was tired. She still tried her best to look after Elwin and me. Of course, I could not have let my mother bear this pressure all on her own, and that is why I would constantly reject my friends’ invitations to whatever activities they were inviting me to join. I had to go home and help my mother as much as possible. Here I have to apologize to my dearest friends. I value friendship, but nothing is more important than my loving mother.

Regardless of how miserable life was, my mother would always care about us first. She is our great mother.

As Mafuyu and Yuki grew, Yuki became more involved in music, playing guitar and starting his band. Mafuyu felt that Yuki was paying less and less attention to him. Deep down, Mafuyu probably hoped that Yuki would invite him to join his band, but he never said that to Yuki. Instead, his loneliness had led to many fights and quarrels between him and Yuki. They were yelling at each other hurtful phrases that were never meant. At one point, Mafuyu said to Yuki, “(If you are willing to do anything for me, then) would you die for me? (じゃ、俺のために死れるの?)”

Undoubtedly, there were times when I would not listen to my mother and yell at her, especially when I was going through my rebellious stage. They were never meant, but hurtful phrases would just come from my mouth. Also, I wouldn’t say I liked it when my mother required me to call her and update her on my whereabouts, such as when I was on a school trip, hanging out with friends, etc. Back then, I felt like having to call my mother in front of my friends was embarrassing, and I probably hoped she would have given me more “trust and freedom” by not having me call her as frequently. I was so wrong that I would have gone back in time and punched myself for having that thought.

Ever since three years ago, I have been reflecting on my whole life as I was watching “GIVEN”. I have never asked my mother if she would die for me, at least not verbally. But actually, I did (kind of) ask her to do that, considering my decision five years ago. To continue pursuing my dream of becoming a physicist, I accepted the offer from Caltech graduate school to continue my PhD studies. I still remember that very night on February 1, 2019, when I received an email from Caltech saying I was accepted. I could not stop myself from screaming for joy. It was so loud that my mother and Elwin were awakened. Elwin told me he thought I was screaming because I saw a mouse. I told them the news, and my mother was extremely happy for me. She hugged me so tightly and tears of happiness flowed down her cheeks.

With that, I left my mother and Elwin behind in Hong Kong as I continued to pursue my dream at Caltech in the US. The selfish me left behind the burden of taking care of Elwin and the family to my mother. Back then, I did “ask” my mother to die for me, unknowingly...

Mafuyu probably said that impulsively without thinking. However, the next day, when Mafuyu came to Yuki's house, he could only find a lifeless Yuki in front of him - Yuki hung himself. Mafuyu stood there forever, unable to utter a single word. He was staring blankly, without any crying. Mafuyu could not express his feelings to anyone, and he decided to lock them away with his true self and continue to live on as if nothing had happened.

January 2020 was the last time I could return to Hong Kong and reunite with my mother. Then COVID-19 happened. Since then, I have only been able to talk to my mother on the phone. Since early April 2021, my mother had not been feeling well, and after my nagging, she finally agreed to go to the hospital and was asked to stay behind because of her high fever. Being in the US, I could only obtain news about her from Elwin, who was also busy with schoolwork as an undergraduate. Knowing that my mother had an abnormally high concentration of white blood cells, all I could do was search the internet with her symptoms for possible illnesses that she could have had - it could have been just a regular fever. Still, it could have also been meningitis, or even blood cancer...The different possibilities have only intensified my feeling of insecurity. My mother was not willing to go to the hospital only until the very end, and I know the reason why. Eight years ago, my mother also went through a severe illness. At that time, I had been nagging her for a long time, asking her to go to the hospital, and she even yelled at me for being annoying; only in the end did she agree to my request. Of course, I fully understand her concerns. My mother cared about Elwin and me and did not want to leave us alone. It was the same this time. I was not home in Hong Kong, so Elwin would have been left alone if my mother had gone to the hospital. She must have been worrying that Elwin would feel scared and lonely, and so she only went to the hospital at the very last moment.

I could fully understand how painful it would have been for Elwin to worry about our mother, and at the same time, needing to take care of everything for the family. That is why I wanted to help reduce Elwin's pressure. Therefore, I called Elwin the second morning after our mother went to the hospital to see if he had any news about her. I requested Elwin to tell me the hospital she was staying at so that I could ask Uncle Eric to check on her. I wanted to have another person to share the worries with Elwin and support the lonely Elwin in Hong Kong. Those were good motives, but I was not smart enough to notice that my constant phone calls to Elwin made

Elwin think I distrusted him. With both of us suffering from extreme worry and stress, we ended up in endless quarrels and yelling. I am grateful that Elwin forgave me for yelling at him and that he eventually understood the motivation behind my actions.

Ever since my mother went to the hospital and fell into a coma, I have been wanting to call her and listen to her voice. But at the same time, I did not want to disturb her from her rest. When I heard from Elwin that my mother's fever had gone down and was able to talk, I could not stop myself from calling her immediately. Even though my mother's voice sounded tired when I heard her saying "hello", "I am really tired" and "I am signing off", I did feel less worried. Hearing my mother's voice, all I could say were "Great", "It is great that you are conscious". Knowing that she was tired, I could only say "rest more" and quickly end the call. If I had known that it was the last time I could talk to her, even though it was very selfish, I would not have ended that call that quickly.

I felt hopeful when I heard my mother was getting better, but Elwin soon told me her situation was worsening again. She fell into a coma again and was sent to the intensive care unit (ICU). My heart sank immediately, but I knew life must go on, so I kept working hard. On 20 April 2021, I confidently presented at the APS April Meeting for the first time as a graduate student. I tried to share my happiness through a Facebook post. People have been asking me why I write long journals on Facebook whenever I go on trips or attain outstanding achievements. In the past, I would answer that it was only personal interest, or maybe it was that I wanted to share my happiness with others, but now that I think about it, I was hoping that my mother could read through my journals and share my joy. I hoped my mother would be proud to have me as her son, and I wanted her to know that her unlimited love for Elwin and me was worthwhile. I was too silly that I was hoping that my Facebook post could have woken her up from her coma. I really hoped that she would be among the many people who liked my Facebook post. That had been my wish.

Who would have thought that I would be receiving the bad news of my mother's death the very afternoon of the same day through a phone call with Elwin? Unlike Mafuyu, my eyes immediately flooded with tears while holding my phone, but I knew I could not cry with Elwin on the line; I could not let Elwin hear me crying because Elwin was already crying loudly as he spoke to me. I knew I had to calm down, and I had to hide my feelings for now; I had to first calm Elwin down, so even though I was in great pain, even though my tears were flowing from my eyes

non-stop, I still forced myself to say “there, there”, “don’t cry”, “Alvin is here, I love you” to the other side of the phone. Only after Elwin felt better and after we ended the call could I truly release my sadness and cry as loudly and as much as I wanted in my dorm room. I must have been crying very loudly because even my neighbors came over to check on me. Eventually, I stopped crying on that very same day. Back then, as I wrote in one of my Facebook posts, I thought, “You cannot cry in front of other people, right? And at the same time, I thought I had cried enough”.

That turned out to be a terrible mistake. I thought that I had cried enough, and similar to Mafuyu, I decided to lock away all my feelings together with my inner true self and forced myself to move on. I acted like nothing happened in front of everyone, especially in front of my friends, colleagues, and loving supervisor, Alan. I put on a bold face, but at the same time, I distanced myself from everyone. I avoided all kinds of interactions and social events with my research group members as much as possible. I did not realize I was doing this on purpose until recently. In front of everyone else, I always put on a big smile and a happy face. I drew various cartoons and designed attractive slides and presentations to make everyone laugh. People probably would just see me as an eccentric, weird graduate student. Probably, that was my purpose - to hide my true self and feelings from everyone, or maybe I wanted to relieve myself from the pain I was suffering by seeing others smile. Back then, I always covered my ears with my earphones, playing music as loud as I could and avoiding all kinds of conversation as much as possible.

┌ Mafuyu isolated himself from everyone else until a year later when he met Ritsuka Uenoyama (上ノ山立夏), another student studying in the same school as him. Uenoyama found Mafuyu lying by the staircase, holding onto the guitar with snapped strings that Yuki had left him with. He helped Mafuyu fix his guitar, eventually invited him to join his band, and taught him how to play the guitar. This eventually brought Mafuyu back into the sunlight, where he also met new friends in the band, although he was still hiding his true feelings from everyone. ┐

I kept isolating myself and let the fake “moved on” me take control of me for more than a year. At some point, my inner self started to resist the control and tried to get out of the prison I threw him into. If my research group members remember, at one group meeting presentation ², I put in a slide to explain that “I am not icy and cold”;

² Actually, I remember the exact date of that group meeting - 22 October, 2022.

that “Although I look brave on the outside, I am actually quiet and shy”, and I even tried to ask people not to “hesitate to say hi to me” because “I will always reply”. Looking back, that was the best attempt of my inner self trying to break through and reach out for help. Fortunately, that attempt worked. After that group meeting, people started to say hi and have daily conversations with me again. People may not know, but even a single hi and simple small talk brighten me and make me feel valued and cared for. Every little act leads to a significant impact. Although, like Mafuyu, I still have not figured out how and when to express my true feelings and who I should talk to. It is not typical to casually bring up the topic that one lost his mother and that one feels sad when talking to your friends and colleagues.

As days passed, Uenoyama turned a melody that Mafuyu hummed into a song and asked Mafuyu to write lyrics, prompting him to fill it with words and thoughts he wanted to convey to the others. Uenoyama hoped to have Mafuyu sing the song in their next public performance. This opened up an opportunity for Mafuyu to face his locked away feelings. It was not easy, especially when he needed to figure out what he wanted to say. In one scene, while another band member, Haruki Nakayama (中山春樹) was asking why Mafuyu found it so difficult to fill in the lyrics for the song, Mafuyu asked the following questions:

“Do you have someone you love, Mr. Haruki? (春樹さんは好きな人っていますか?)”

“If that someone suddenly disappeared from the world one day, what would you say? (もし、その人がある日突然世界からいなくなったら、なんて言います?)”

“Would you feel lonely? Sad? Would you miss them? (悲しい? 寂しい? 恋しい?)”

Obviously, Haruki was not able to answer these hard questions, and he could only reply that “I honestly don’t know what to say...” (何て言えばいいか分からないかな...), which Mafuyu agreed to.

I did not have such an opportunity as Mafuyu, but ever since I heard those questions and words from Mafuyu, I cannot stop thinking about them. Figuring out the answers to the questions is hard. It is not like I am trying to escape from facing the real me, but I am genuinely confused and I do not know what my true feelings are - Too many mixed emotions. Do I feel lonely? Do I feel sad? Do I miss my mother?

Do I regret leaving my family behind? Do I feel angry towards myself? Do I feel angry towards my mother for leaving me and Elwin so suddenly? I do not know the answers, nor do I know how I should express them. It is too complicated.

「Mafuyu kept pondering the questions day after day, and even until they had their final rehearsal before the performance, he failed to come up with the lyrics for the song. But just before the performance, Mafuyu heard his inner voice talking to him: “I never know what kind of reactions I should be giving. I am not able to laugh and cry normally like everyone else does. It is probably because I am bad at expressing myself to the others. But even so, in reality, even if I really want to, I cannot cry properly, and it hurts so bad. It hurts so much that I just want to scream. I always wanted someone to understand that, even just a little bit...” (いつもどういう反応したらいいかわからないんだ。他の人がするみたいに、泣いたり笑ったり、上手に出来ない。多分きっと人より伝えるのが下手くそなんだ。でも、だけど本当は、いつも本当は、ただ、上手に泣けなくて、苦しくて、痛くて、叫びたいのは、誰かにずっと分かって欲しかった。少しだけでいいから...) And with that, Mafuyu finally expressed his true feelings to everyone through the song “A Winter story” (冬のはなし[175])...

I think I am the same as Mafuyu. I am terrible at expressing myself to others, and I cannot laugh and cry (genuinely) in front of others. And because I cannot express myself properly, it hurts so bad. Tears will start flowing endlessly down my cheeks whenever I am alone, at night in my bedroom or in the office. Mafuyu’s “A Winter Story” resonates with my feelings so much that I cannot stop myself from looping the song repeatedly, and I am pretty sure that I have listened to that song not fewer than ten thousand times over the past few years. If you have ever stumbled upon me strolling on the campus wearing my earphones, that is actually your closest approach (periapsis) to my true self, when listening to Mafuyu’s singing, as if he was saying exactly what I wanted to say. I try my best to hide my inner self with a big smile whenever I see someone I know, but if you look closely, you will probably notice my watery eyes. For completeness, I include the lyrics of the song below.

A Winter Story (partial)

冬のはなし（部分）

また 溶けきれずに残った
日陰の雪みたいな
想いを抱いて生きてる

I continue with feelings inside me that is
not disappearing, like the snow that has
not completely melted in the shade.

ねえ 僕はこの恋を
どんな言葉で
閉じたらいいの？

Please tell me what words I should use to
close the door on this only love I know?

あなたの全てが
明日を無くして
永遠の中を彷徨っているよ
さようなら出来ずに
立ち止まったままの
僕と一緒に

Your everything has lost its tomorrow,
and it is now wandering around eter-
nally. Alongside me, who was unable
to say goodbye to you or move on.

また 解けない魔法のような
それとも呪いのような
重い荷物を抱えてる

Just like a spell that will not break, or like
some kind of curse, I am still holding
onto some heavy luggage that you left
behind.

ねえ 僕はこの町で
どんな明日を
探せばいいの？
嗚呼！

Please tell me what kind of tomorrow
am I supposed to look for in this world?
(Scream)

冷たい涙が
空で凍てついて
優しいフリして
舞い落ちる頃に

The cold tears that fall freeze under the
sky. They pretend to be kind as they flow
down my cheeks as if they are dancing.

離れた誰かと
誰かがいたこと
ただそれだけの話

Two people who were always together
were torn apart. That is all to the simple
story I am telling you.

The lyrics of Mafuyu's song is the same message I want to convey. I lied to everyone,
even myself, that I was feeling okay. In reality, I have never moved on since my

mother passed away three years ago. I have been standing in the same place, unable to carry on. I tried to hide away all my feelings and my true self, but they continued to live with me like the snow that was not melting in the shade. I cannot find the right words to properly close the door on the only love I have ever known from my mother, nor to express my feelings. All those memories and emotions that my mother left me with are like a spell or a curse that cannot be broken. In the end, the thing that causes me the most pain, which I will regret for the rest of my life without any solution, is that I can never properly say goodbye to my mother. Without her, I am not even sure what kind of tomorrow I should be looking for in this world.

As Mafuyu was singing, his inner voice spoke again: “This is a story that happens in a winter. This is a story that happens over a night. Wherever I go, you are always here. You are always living inside my head. You are living in all the sceneries I am seeing. Wherever I go, there are always memories of you lingering. I cannot forgive you, nor can I forgive myself. But I wish I can forgive both of us. I am really lonely.” (とある冬の話。とある夜の話。どこに行っても君がいるんだ。俺の頭の中に君がいるんだ。景色の中に君がいるんだ。どこに行っても、君がいた匂いと景色があるんだ。君を許せない。俺を許せない。許したい。寂しいよ。)

It is the same for me. I tried my best to forget about the sadness, but ironically, I also do not want to forget my mother and all the memories, happy or sad, I have with her. Wherever I go, I see shadows of her. When I see children playing happily with their parents on the campus over the weekend, or when it comes to all those holidays when people would go home and reunite with their friends and relatives, I would feel sad and jealous, staying behind all alone and depressed. It is not very physicist-like to say this, but I can feel my mother still being together with me. She is always there wherever I go. I still cannot forgive her for leaving me. I also still cannot forgive myself for leaving her in the first place. But I do wish to forgive both of us and let go. I feel really lonely. How I wish I could talk to you again and hear you telling me that I did a great job and that you are proud of me. But that will never happen...

As everyone continues to talk to me and brings me out into the sunlight, my inner self also tries harder and harder to take back control from the fake me. The pain and the hope to find someone to understand me continue to grow daily. I know I cannot hold my inner self in his prison much longer. And therefore, I decided I needed to

find a time to end this properly. That is why I really hope to have my thesis defense on the week of April 25, 2024. That week marks the third year since my mother passed away. It would also have been her birthday if she was still alive. I want to use the passing of my thesis defense as a final present to my mother, and as I finish my Ph.D. studies, I will also use it to finally say a proper goodbye to my mother and end one of the most important chapters of my life. I wish I could do that... (Well...turns out I have to defend my thesis a week later, but as Mafuyu says, “it is going to be fine (大丈夫。)”.)

And that is the simple story I want to tell everyone reading my thesis. Apologies for the lengthy passage, but I really need to use this opportunity to express my true feelings and explain why I would like to dedicate my thesis to the manga and anime series “GIVEN”. It helped me to reflect upon myself, release my true self, face my hidden feelings, and get out of the darkness into the sunlight again. It has been an essential part of my PhD journey.

For the rest of my acknowledgment, I would like to thank many people properly.

First and foremost, I would like to thank my most respected teacher and supervisor, Professor Alan J. Weinstein. Alan has been the most loving, caring, and supportive teacher throughout my PhD journey. He gives me the most freedom to work on things that interest me, and he will never force projects or work on me. He provides the best guidance and suggestions to advance my research path. At times, he would remind me to expand my expertise and not only focus on one single research topic. He introduced me to postdocs (e.g., Dr. Ryan Magee) with similar research interests to mine and provided alternative research opportunities for me. Alan is also a very attentive and kind-hearted supervisor. He can catch the slightest sadness and bitterness within my conversation with him, and he always finds time to talk to me, comfort me, and provide the best solutions for whatever problems or concerns I have. I will never forget the four years of back and forth with the reviewers regarding my first academic paper. I had been having a hard time with that, and Alan has always been standing next to me, supporting me with his kind words and helping me publish my paper. I still remember the day after my paper was accepted, I saw Alan on the second floor of West Bridge, and he gave me a pat on the back and said “Good job” to me. I still remember Alan coming out of the office and giving me a high five when he saw the cartoon of himself I drew during my first LIGO SURF lecture on Lensing back in 2021. I still remember Alan spending much time helping me figure out my future path. I still remember Alan cheering me up and

telling me how to appreciate the research and work I have been doing and not to diminish myself...Alan means a lot to me, especially to international students like me who study all by themselves far away from home. Alan treats his students and group members as part of his family, and you can see that when he always uses “hey kids” to say hello to us. Seeing and talking to Alan makes me feel warm and loved. At some point in my PhD journey, Alan has already become like a caring and loving father to me, and I feel safe to share everything with him. In fact, Alan is the only one in the group (apart from Dr. Rico Lo) to whom I have told about my mother. I appreciate and will forever be grateful for everything Alan has done for my brother Elwin and me, all the support he has given me, and all the time he spent with me throughout my SURF and PhD journey. Having Alan as my PhD supervisor, mentor, and teacher is one of the best things that has ever happened in my life. Thank you very much, Alan!

I would also like to thank Juno Chan. We first met as mentor and mentee in 2019, but since then, you have become a close collaborator of mine and my dearest friend. Together, we built the TESLA search pipeline, which remains LVK’s flagship pipeline for searching for possible sub-threshold lensed gravitational waves. I will never forget when we “fought” all the “academic battles” together. You are always here by my side to deal with all my research work difficulties. Apart from work, you are always willing to spend time to chat with me. My PhD journey has always been challenging, especially alone in the US. You have been a good listener and a great friend who always offers me constructive advice. Whether I am sad or angry, you are here to lend me an ear, listen to me, and provide me with all the support you can or ways to view things from a different angle. You always help me to locate the root of all my negative emotions. You are always one of the sources of happiness for me. Simply talking to you can already relieve a ton of pressure and sadness. You are my dearest friend, not just because you always stand by me but also because you are willing to scold me when I am not on the right track. Sometimes, I would be blinded by anger or emotions and cannot make decisions logically. You are always be there to yell at me and help me get back to my normal state. Without your unlimited support, I cannot imagine how I would have turned out. Thank you, Juno, for being my dearest friend, and I hope we will continue to be dearest friends forever. I wish you all the best in your PhD journey and future. We will continue to support each other!

Next, I would like to thank Dr. Heather Fong. Since the early stages of developing

TESLA, you have provided all the help and support we needed regarding reviews and development projects. You are always willing to listen to all our difficulties in the LVK lensing community or the GstLAL world and provide concrete and helpful advice and solutions for us to move forward. Although we had never met in person until the LVK September meeting in 2023, you have always been a great friend of Juno and me. You are always willing to find time to hear about our problems with the projects we are working on and with other collaborators. Outside of work, you are also a wonderful friend. The time we spent touring around Japan was one of the happiest times I had throughout my PhD journey. It was also when I smiled and laughed genuinely because of all our fun. They are some of the most incredible memories that I will treasure forever. I will always be grateful for all your help and support, Heather. I wish you, your mother, and your family all the best in the future. Please come and visit me in Japan! ありがとうございます！

I want to thank Dr. Rico Lo, my mentor, friend, and roommate. I have learned many things from you regarding how to do research. As a friend and roommate, your emotional support was critical in keeping me from breaking down throughout my PhD journey. Thank you very much for everything.

I want to thank my Caltech LIGO laboratory friends and colleagues (current and past), including Professor Katerina Chatziioannou, Dr. Derek Davis, Dr. Virginia d’Emilio, Dr. Lucy Thomas, Dr. Colm Talbot, Dr. Arianna Renzini, Dr. Jonah Kanner, Dr. Patrick Meyers, Sophie Hourihane, Isaac Legred, Simona Miller, Brian Seymour, Shreya Anand, Taylor Knapp, Rhiannon Udall, Jacob Golomb, Ethan Payne, and Shruti Maliakal, for all their love, help and support. Thank you for giving me your unlimited support and love. It is really a blessing to know everyone here and to be part of the Caltech LIGO family. As I said before, all your small talk, or even just a simple “hi”, has made a big difference in my life.

I want to thank Professor Tjonnie Li for being my first supervisor and introducing me to the world of gravitational-wave physics research. Without your constant support, both work-related and emotional, I would never have made it this far. It is a blessing to have you as one of my supervisors in my academic career, and I will be forever grateful for everything you have taught and done for me. Thank you, Tjonnie!

I want to thank Professor Ming Chung Chu. Professor Chu was one of the interviewers when I was at the admission interview for CUHK’s Enrichment Stream for Theoretical Physics undergraduate program. I mentioned that I had never joined any physics-related competitions or contests and asked Professor Chu if that meant

I would not have a chance to be admitted to the program. Professor Chu's reply strengthened my confidence in pursuing physics at the university, and throughout my academic career, he has also been supportive and helpful.

I want to thank Professor Kipp Cannon for his guidance and help, and for providing me the opportunity to learn from and work with him as a postdoctoral fellow.

I want to thank Dr. Ryan Magee for his mentoring in several GstLAL-related projects and for suggesting possible career paths.

I want to thank Professor Surabhi Sachdev for being my LIGO SURF mentor in 2018, for introducing me to the world of research, and for being a great friend of mine.

I want to thank Professor Leo Tsukada for being a great friend and colleague of mine.

I want to thank Richard George for being a great friend and colleague of mine ever since we met in 2022.

I want to thank my high school physics teacher, Mr. Chi Leung Wong, for inspiring me with advanced physics and nurturing my interest in pursuing physics.

I want to thank Dr. Ken Ng for his support and help throughout my PhD journey.

I want to thank all my SURF students, including Darin Mumma, Storm Colloms, Aidan Chong, and Haley Boswell, for all their hard work and passion for doing research in gravitational-wave physics. Working with you all every year helps me strengthen my love for LIGO and gravitational waves. I wish you all the best in your future.

I want to thank my friends and relatives for their unlimited love and support.

I want to thank all my teachers for their teaching and care.

I want to thank my dearest brother, Elwin Li, for being a most supportive brother. Elwin is the earliest student of mine, whom I have gradually led to become interested in physics. I am genuinely grateful that he has forgiven me for leaving him and my mother behind in Hong Kong and pursuing my dream of studying physics at Caltech. I will be forever proud to have you as my brother, and you are always the best Christmas present I can ever get. I wish you all the best in the future!

Last but not least, I would like to express my deepest, sincerest thanks to my mother, Wendy Choy, for sacrificing her whole life to nurture Elwin and me. You are, and

will forever be, the most incredible mother in the world. Although I can never speak to you again, I am confident I will live the best of my life and make you as proud as possible to have me as your son. You will live forever in my heart. May you rest in peace.

I apologize again for the long acknowledgment, for I am grateful to everyone who has once appeared in my life. Without any one of you, there will be no Alvin today.

ABSTRACT

The LIGO-Virgo-KAGRA (LVK) collaboration first observed gravitational waves in 2015, and more than 90 gravitational-wave events have been observed, all coming from mergers of compact objects (black holes and neutron stars), known as compact binary coalescences (CBC). Studying and observing gravitational waves opens a new window for us to understand the nature of spacetime and the universe. Strain data from LVK's detectors are analyzed by search pipelines to identify weak gravitational-wave signals in noisy data. To maximize the potential of gravitational waves, it is essential to continue to improve search pipelines' sensitivity to probe GW sources with the broadest range of parameters and from the furthest distances. I will give a detailed overview of the GstLAL pipeline and present related development (ongoing) work for GstLAL to enhance its search effectiveness and efficiency.

In the second part of my thesis, I will focus on gravitational lensing of gravitational waves. As masses can produce curvature in spacetime, gravitational waves, like electromagnetic (EM) waves, are deflected when passing by massive intervening objects before reaching gravitational-wave detectors on Earth, an effect known as gravitational lensing. Observing lensed gravitational waves confirms another prediction in Einstein's general relativity and enables us to conduct cosmography studies, test general relativity, search for dark matter and other exotic phenomena, and deepen our understanding of the universe. I will give a detailed introduction to gravitational lensing of gravitational waves. We then introduce a Targeted sub-threshold search for strongly-lensed gravitational wave pipeline called "TESLA". The TESLA pipeline is the flagship to look for sub-threshold lensed gravitational waves. Next, we present the results of the LVK collaboration-wide effort to search for lensing signatures in gravitational-wave data from the third observing run O3. Next, we introduce a significant update to the TESLA pipeline, now known as the TESLA-X pipeline, with enhanced search sensitivity towards lensed gravitational waves. We also introduce an alternative ranking statistic implemented into the TESLA-X pipeline that considers the signal's consistency with the assumed lens model. Finally, we end the thesis with a summary and an outline of possible future work.

PUBLISHED CONTENT AND CONTRIBUTIONS

- [1] **Alvin K. Y. Li** et al., “Targeted subthreshold search for strongly lensed gravitational-wave events”, *Phys. Rev. D* 107.12 (2023), DOI: 10.1103/PhysRevD.107.123014. (**Alvin K. Y. Li** is the first author of this paper. He is the main developer of the “TargetEd Subthreshold Lensing seArch” (TESLA) method to find possible sub-threshold lensed counterparts to known gravitational waves described in this manuscript and led the writing of the manuscript.)
- [2] **Alvin K. Y. Li** et al., “TESLA-X: An effective method to search for sub-threshold lensed gravitational waves with a targeted population model”, *arXiv e-prints*, 2023, DOI: arXiv.2311.06416. (**Alvin K. Y. Li** is the lead author of this paper, which describes the next generation “TargetEd Subthreshold Lensing seArch” (TESLA) method with the use of better-constructed reduced template bank and a targeted population model.)
- [3] Y. Wang, R. K. L. Lo, **Alvin K. Y. Li** and Y. B. Chen, “Identifying Type II Strongly Lensed Gravitational-Wave Images in Third-Generation Gravitational-Wave Detectors”. *Phys. Rev. D* 103.10 (2021). DOI: 10.1103/PhysRevD.103.104055. (**Alvin K. Y. Li** led the mock simulation campaign to search for type II lensed gravitational waves with a type I image template bank using TESLA, as described in section III in the paper. He also edited Section III of the paper.)
- [4] LIGO Scientific and Virgo Collaborations, R. Abbott, ...(723 authors)..., **Alvin K. Y. Li**, ...(655 authors)..., “Search for Lensing Signatures in the Gravitational-Wave Observations from the First Half of LIGO–Virgo’s Third Observing Run”, *Astrophys. J.* 923.1 (2021), DOI: 10.3847/1538-4357/ac23db. (**Alvin K. Y. Li** led the analysis for the search for sub-threshold lensed gravitational waves and edited the corresponding section in the LIGO-Virgo-KAGRA collaboration-wide paper (Section 5.3).)
- [5] LIGO and Virgo Scientific Collaboration, R. Abbott, ...(850 authors)..., **Alvin K. Y. Li**, ...(842 authors)..., “Search for gravitational-lensing signatures in the full third observing run of the LIGO-Virgo network”, *arXiv e-prints*, 2023. DOI:10.48550/arXiv.2304.08393. (**Alvin K. Y. Li** is one of the editorial leads for writing this LIGO-Virgo-KAGRA collaboration-wide paper and led the analysis for the search for sub-threshold lensed gravitational waves (Section 3.1).)
- [6] J. Janquart, ...(6 authors)..., **Alvin K. Y. Li**, ...(22 authors)..., “Follow-up Analyses to the O3 LIGO-Virgo-KAGRA Lensing Searches”, *Mon. Not. Roy. Astron. Soc.* 526.3 (2023), DOI:10.1093/mnras/stad2909.

(**Alvin K. Y. Li** led the analysis and provided results for the search for GstLAL-based sub-threshold lensed gravitational waves described in the paper. He also helped with editing the corresponding sections in the paper.)

- [7] S. Goyal, S. Kapadia, J. R. Cudell, **Alvin K. Y. Li**, and J. C. L. Chan, “A rapid method for preliminary identification of subthreshold strongly lensed counterparts to superthreshold gravitational-wave events”, *Phys. Rev. D* 109.2 (2024), DOI:10.1103/PhysRevD.109.023028.

(**Alvin K. Y. Li** led the analysis and provided results for the search for GstLAL-based sub-threshold lensed gravitational waves for follow-up analysis, as described in the paper. He also helped with editing the corresponding sections in the paper.)

- [8] B. Ewing, ...(7 authors)..., **Alvin K. Y. Li**, ...(35 authors)..., “Performance of the low-latency GstLAL inspiral search towards LIGO, Virgo, and KAGRA’s fourth observing run”, *arXiv e-prints*, 2023, DOI:10.48550/arXiv.2305.05625.

(**Alvin K. Y. Li** helped to develop, operate, and monitor the mock data challenge with GstLAL as described in the paper.)

- [9] **Alvin K. Y. Li** et al., “TESLA-X+: Targeted search method for sub-threshold strongly lensed gravitational waves with lens-model-based ranking statistics”, In Preparation.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	xx
Published Content and Contributions	xxi
Table of Contents	xxiii
List of Illustrations	xxvi
List of Tables	xlvi

Part I: Introduction to gravitational waves and gravitational-wave data analysis with the GstLAL Search pipeline . . .

Chapter I: Overview of the thesis	2
Chapter II: Introduction to gravitational waves	4
2.1 Overview of this Chapter	4
2.2 A very brief overview of General Relativity	4
2.3 Linearized gravity in the weak-field limit	5
2.4 Gravitational waves	13
2.5 Energy carried by gravitational waves	21
2.6 The LIGO detectors: An overview	25
Chapter III: Data analysis of gravitational waves from compact binary coalescences - An overview	32
3.1 Introduction	32
3.2 Overview of current LVK's framework to search for gravitational waves from compact binary coalescences	32
Chapter IV: The GstLAL search pipeline - An overview	40
4.1 Top-level introduction to matched-filtering-based search pipelines	40
4.2 An overview of the GstLAL search pipeline	43
Appendices	61
4.A Deriving the normalization constant for GstLAL's auto-correlation based signal consistency test value ξ^2	61
Chapter V: Improving the ranking statistic assignment of the GstLAL search pipeline with the inclusion of iDQ data-quality information	65
5.1 Introduction and Overview	65
5.2 Background	65
5.3 Proposed method	68
5.4 Expected results	70

Part II: Gravitational Lensing of Gravitational Waves	72
Chapter VI: Introduction to gravitational lensing and lensing of gravitational waves	73
6.1 Overview	73
6.2 A very brief overview of gravitational lensing	73
6.3 Gravitational lensing of gravitational waves	79
6.4 Overview of the upcoming chapters	89
Appendices	91
6.A Gravitational lensing in Newtonian Mechanics	91
6.B Derivation of the deflection angle in Newtonian Mechanics	92
6.C Relativistic deflection angle	101
6.D The lens equation from geometric constructions	106
6.E Deriving the time delay components induced by gravitational lensing	113
Chapter VII: The Targeted subthreshold search for strongly lensed GWs (TESLA)	116
7.1 Introduction	116
7.2 The TESLA search method for subthreshold lensed gravitational waves	117
7.3 Simulation campaign	120
7.4 Concluding Remarks	133
Chapter VIII: Search for lensing signatures in the first half of LVK's third observing run	136
8.1 Introduction	136
8.2 Data and events considered	137
8.3 Lensing statistics	139
8.4 Analyzing high-mass events	145
8.5 Search for multiple images	148
8.6 Search for microlensing effects	167
8.7 Conclusions and outlook	170
Appendices	172
8.A Lensing statistics supplementary	172
8.B Construction of sub-threshold counterpart search template banks	173
8.C Injection study for microlensing analysis	175
Chapter IX: Search for lensing signatures in LVK's full third observing run	176
9.1 Introduction	176
9.2 Data and Events	177
9.3 Strong lensing	179
9.4 Microlensing Effects	189
9.5 Implications	190
9.6 Concluding Remarks	198
Appendices	200
9.A Type II simulation campaigns	200
9.B Marginalized Posteriors of Microlensing Parameters	202
Chapter X: Follow-up Analyses to the O3 LIGO-Virgo-KAGRA Lensing Searches	206
10.1 Introduction	206

10.2 Analysis Methods	208
10.3 GW190412	211
10.4 GW191103–GW191105	214
10.5 GW191230_180458–LGW200104_180425	227
10.6 GW200208_130117	234
10.7 Conclusions and Prospects	241
Appendices	258
10.ALensID GW191103–GW191105 Investigations	258
10.BDiscarded Targeted PyCBC Sub-Threshold Search Triggers	259
10.CDetails of the Millilensing Investigation	259
Chapter XI: TESLA-X: A more effective targeted search method for sub-threshold strongly lensed gravitational waves	265
11.1 Introduction	265
11.2 The search for sub-threshold strongly-lensed gravitational waves: Objectives and current approach	266
11.3 The upgraded TESLA-X search method	269
11.4 Method verification and results	276
11.5 Concluding Remarks	282
Appendices	285
11.AChoosing a suitable pair of parameters to construct the Gaussian KDE	285
Chapter XII: Re-ranking possible sub-threshold lensed candidates with lensing likelihoods	295
12.1 Introduction	295
12.2 Motivation and Background	296
12.3 The Lensing Likelihood	299
12.4 Working example	301
12.5 Conclusion and future work	303
Part III: Summary	305
Chapter XIII: Summary and future work	306
13.1 Summary of Thesis	306
13.2 Future work	308
Bibliography	311

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.4.1 How a ring of test masses on the x - y plane deforms as a gravitational wave with plus-polarization travelling in the z -direction passes by.	21
2.4.2 How a ring of test masses on the x - y plane deforms as a gravitational wave with cross-polarization traveling in the z -direction passes by.	21
2.6.1 A simplified schematic diagram of the Advanced LIGO gravitational-wave detector. Figure from [341].	25
2.6.2 A schematic diagram illustrating gravitational-wave source frame and detector frame.	30
3.2.1 A block diagram of the feedback control loops used in the LIGO detectors. Figure taken from [341].	33
3.2.2 An illustration of how photon calibrators imping laser beams on a suspended test mass. Figure taken from [207].	34
3.2.3 The 15 parameters that govern the waveforms of gravitational waves from compact binary coalescences.	36
4.2.1 A simple flowchart outlining the key components of the GstLAL pipeline. Figure reproduced with reference to [250]. Note that the first two boxes representing “data and template preparation” and “data filtering” are done individually for each participating detector.	43
4.2.2 The median power spectral density (PSD) of data collected from the Hanford (H1, in red) and Livingston (L1, in blue) detector between UTC time 2023-05-24 15:00:00 and 2023-06-06 14:30:34, estimated using the GstLAL pipeline. The binary-neutron-star horizon distances for the two detectors are also estimated and shown in the legend. Note that GstLAL downsamples the data being analyzed, and hence the PSD shown above ends at 10^3 Hz. However, the PSD extends beyond 10^3 Hz in reality.	45

- 4.2.3 The template bank used by the GstLAL search pipeline in O3, shown in the component masses $m_1 - m_2$ space. The full bank is composed of five smaller banks, including the binary-black-hole (BBH) bank (in blue), binary-neutron-star (BNS) bank (in orange), neutron-star-black-hole (NSBH) bank (in pink), intermediate-mass-black-hole (IMBH) bank (in green) and other BBH bank (in red). These different banks are used to target different regions in signal subspace. 46
- 4.2.4 An example of two split banks (bank 1000 and 1019) generated from the full template bank, shown in the chirp mass \mathcal{M}_c - effective spin χ_{eff} space. Note that the partially overlapping region exists only for intermediate use when adopting the LLOID method, but will be removed afterward so that there are no repeated templates. 47
- 4.2.5 An example of the time slices for a template waveform with chirp mass $\mathcal{M}_c = 2.61 M_\odot$ and z-component spins $s_{1z} = -0.73$ and $s_{2z} = -1.0$. Note that the sampling frequency increases with time as the Nyquist frequency increases. 48
- 4.2.6 A schematic diagram of the singular-value-decomposition (SVD) operation on the template bank. 49
- 4.2.7 The auto-correlation timeseries (blue striped line) and SNR timeseries (black solid line) evaluated for GW230529 for the LIGO Livingston (L1) detector. In the above figure, the trigger's coalescence time is set to $t = 0$ 50
- 4.2.8 Triggers found in time coincidence (within a certain time window that accounts for (1) the maximum travel time of gravitational waves between detectors, and (2) statistical fluctuations in the measured trigger time due to detector noise) between participating detectors with the same template parameters are grouped and labelled as coincident (signal) triggers by GstLAL. Coincident triggers are further followed up by the pipeline as signal triggers. In contrast, single detector triggers (non-coincident triggers) are treated as noise (background) triggers by GstLAL, and are used to construct the noise background. Starting from LVK's second observing run O2, GstLAL also allows single detector triggers to be treated as signal candidates if they pass a preliminary log-likelihood ratio threshold. This will be further explained in later subsections. 52

- 4.2.9 The ρ - ξ^2 histogram for noise triggers collected in the Hanford (H1) detector for one split bank between GPS time 1370097052 and 1370356252 (3 days of data). The green circle marks an example signal trigger with H1 SNR $\rho_{\text{H1}} = 10$ and $\xi_{\text{H1}}^2 = 1$. Since it is (marginally) away from the background noise triggers distribution, GstLAL will evaluate a low $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ for this trigger, in turn increasing the log-likelihood ratio and hence the significance of this trigger. 54
- 4.2.10 (Left) At different sky locations, a gravitational wave may reach one detector before reaching the other. The effective strength of the signal may also differ. Depending on the ratio of the signal strengths (effective distance ratios), the expected time and phase delay distribution for a pair of detectors has a non-trivial probability distribution, given that they are seeing the same signal. (Right) A cartoon illustrating an example of the Δt , $\Delta\phi$ probability distribution for a trigger seen in both the Hanford (H1) and Livingston (L1) detector with an identical SNR of 8. Assuming both detectors have similar sensitivity, it is more likely for the signal to be detected at the same time with almost no difference in the coalescence phase. Note that this is just a cartoon illustration. See Figure 7 of [309] for a realistic example. 55
- 5.2.1 An example timeseries of log-likelihood estimate $\ln \mathcal{L}_{\text{iDQ}}$ of the presence of non-Gaussian noise in the data output from the iDQ pipeline. The timeseries covers roughly 5000s of data (i.e. ≈ 1.3 hours). Note that $\ln \mathcal{L}_{\text{iDQ}}$ has a value of ≈ 9 most of time, but occasionally it can increase to 70 or more (e.g. see the peak at GPS times ≈ 2600 s). These peaks represent times the iDQ pipeline believes glitches are present in the data. 68

- 5.3.1 The proposed scheme on modifying the original $P(\rho, \xi^2|\text{noise})$ term in GstLAL's likelihood ratio calculations. We propose promoting the 2D ρ - ξ^2 histogram for noise triggers to 3D histograms such that the renormalized iDQ log-likelihood $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ as an additional parameter. In the figure, the proposed 3D histogram is represented by a series of slices of 2D ρ - ξ^2 histograms along the iDQ dimension. GstLAL will then populate the extended 3D histograms for each detector with noise triggers, according to the data quality around the time of each trigger informed by $\ln \hat{\mathcal{L}}_{\text{iDQ}}$. Then, as usual, the histograms will be normalized and smoothed with a Gaussian smoothing kernel. The end product will be a three-dimensional probability density function $P(\rho, \xi^2, \ln \hat{\mathcal{L}}_{\text{iDQ}}|\text{noise})$, which can be used to replace the original $P(\rho, \xi^2|\text{noise})$ term in the likelihood ratio calculation in GstLAL. . . 69

- 5.4.1 An illustration showing how the proposed method can improve the search sensitivity of GstLAL. Suppose we have a signal trigger found in a detector when the data quality is good. A green star in the figure represents the signal trigger. In the current GstLAL framework, since data quality information is not used, the 3D ρ - ξ^2 - $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ histograms (middle) for noise triggers are essentially “marginalized” over the data quality (iDQ) dimension, resulting in the usual 2D ρ - ξ^2 histogram on the left. As discussed before, non-Gaussianity (e.g. glitches) in the data can pollute the histograms and effectively broaden the probability distribution $P(\rho, \xi^2 | \text{noise})$. As shown in the figure, the trigger found during times with good data quality lies in the region of the marginalized probability distribution $P(\rho, \xi^2 | \text{noise})$ with high probability density, that is polluted by noise triggers associated with non-Gaussian noise in the data. Essentially, this means that $P(\rho, \xi^2 | \text{noise})$ for this trigger will be unnecessarily high, resulting in a lower log-likelihood ratio and significance being assigned to it. In the proposed method, on the other hand, the trigger found in times with good data quality effectively has its log-likelihood ratios assigned using a histogram that suits the data quality around the time of the trigger (e.g. right upper histogram). As we can see, the trigger does not live within the probability distribution $P(\rho, \xi^2 | \text{noise})$ for noise triggers coming from times with good data quality, and it will be assigned a lower $P(\rho, \xi^2 | \text{noise})$ and hence higher log-likelihood ratio and significance. This helps to prevent unnecessary down-ranking of triggers, which may lead to potential loss of gravitational-wave signals. 70
- 6.2.1 A cartoon showing a typical case of optical lensing with a convex lens. Light from the object on the left is being focused onto the observer on the right by the lens. The observer does not know that the light has been refracted and will think that it originates from an image (in pink) at a further position (different from the object’s position), with a larger size (magnified). 74

6.2.2	Parameters for deriving the lens equation. Here, the solid black line joining the observer and the center of the lens defined the z -axis. All angles described in this figure are two-dimensional azimuthal angles (e.g. one can decompose the angles into x and y components). $\vec{\theta}$ is the 2D angle of the image position. $\vec{\xi}$ is the 2D impact parameter. $\vec{\eta}$ is the 2D source position. $\vec{\beta}$ is the 2D angle of the source position. All other parameters have the same meanings as in Figure 6.D.2.	76
6.2.3	Critical curves and caustics of a point mass lens model.	79
6.3.1	(Left) Type 1 (blue) and Type 2 (orange) gravitationally-lensed gravitational waveforms in the time domain. (Right) Type 1 (blue) and Type 3 (orange) gravitationally-lensed gravitational waveforms in the time domain. Type 1 signals have zero Morse phase shift, and hence, they are in phase with the not-lensed gravitational wave. Type 2 signals have a $\frac{\pi}{2}$ Morse phase shift compared to the not-lensed gravitational wave. Type 3 signals have a π Morse phase shift compared to the not-lensed gravitational wave.	86

- 6.3.2 The general framework of the LIGO-Virgo-KAGRA (LVK) collaboration's lensing searches. We start with a set of confirmed gravitational-wave events. We perform microlensing (via the pipeline GRAVELAMPS) and millilensing analysis for each event. We also search for possible sub-threshold lensed counterparts to each event via two independent search methods: (1) GstLAL-based search pipeline TESLA, and (2) PyCBC-based sub-threshold search. From the final pool of gravitational-wave candidates, we form pairs of candidates. For each pair, we utilize three independent preliminary tests/pipelines to check for signal consistency, namely (1) posterior overlap (checking for consistency in signal source parameters), (2) machine-learning-based classification test LensID (checking for consistency in sky location of source and time-frequency maps of signals), and (3) phazap (checking for consistency in the signal phase with the strong lensing hypothesis). Each of these tests assigns a ranking statistic to the pairs of candidates, and those that satisfy a pre-determined threshold are passed on for two independent, robust tests to determine how likely they are strongly-lensed counterparts to each other from the same source, namely (1) Golum and (2) Hanabi. Pairs that survive the tests with strong evidence of being strongly lensed will then be followed up further. 88
- 6.A.1 (Left) Objects projected from the Earth's surface can have very different trajectories depending on the initial conditions. (Right) Some objects in the universe may have accumulated sufficiently high masses. Gravity is so strong that not even light can escape. 91
- 6.A.2 Light coming from a star-forming galaxy passes through a galaxy cluster that acts as a gravitational lens to deflect the light rays. Observers on the other side of the lens will think that the deflected light comes from images at positions different from the source position. The images may be magnified / de-magnified in terms of brightness and size, and may suffer from shape deformation. 93
- 6.B.1 An astronaut is riding in a rocket originally moving at speed v horizontally to the right. A blue giant of mass M nearby acts as the gravitational lens. The *half*-deflection angle is $\hat{\alpha}$, and the distance d is the impact parameter. Inside the rocket, the astronaut shoots a light beam horizontally (in red). 94

6.B.2	Here, we show a rocket (the rectangle box) moving horizontally along the positive x -axis initially. A blue giant of mass M (in blue) is placed at the origin. \vec{r} denotes the vector pointing from the blue giant to the rocket, and θ is the angle between \vec{r} and the positive y -axis. $\hat{\alpha}$ is the deflection angle, and d is the impact parameter. A light beam is shot horizontally inside the rocket (not shown).	95
6.B.3	A gravitational lens of mass M is placed at the origin. The x -axis is aligned such that the positive- x direction is along the direction of the un-deflected light ray from the source. The distance from the lens to the closest approach of the light ray is ξ , at a polar angle θ_0 . The position of the light ray at any time t is represented by the position vector \vec{r} . The half-deflection angle is $\delta\theta$, so the total deflection angle is $2\delta\theta$ (deflection angle is defined as the angle between the image and the source).	97
6.B.4	Additional geometric construction on 6.B.3.	102
6.D.1	The “simplest” lens model is the case when no gravitational lens exists. In this scenario, however, we can already define two important quantities, namely the source position β and source distance D_S , which will be useful later.	107
6.D.2	A massive point-mass lens enters the picture. Because of gravitational lensing, the light ray from the source is bent, and a new image is formed. We now need to define several new parameters: (1) The image position θ , (2) The lens-source distance D_{LS} , (3) The lens distance D_L , (4) The impact parameter ξ (the distance of closest approach between the light ray and the lens), and (5) The deflection angle $\hat{\alpha}$ (the angle between the un-deflected light ray and the image).	108
6.D.3	The problem of solving the image positions is transformed into a graphical problem, asking for the number of intersection points between two curves $y = \theta - \beta$ and $y = \frac{\theta^2}{\theta_E}$. As we can see, the point mass lens model always has two images (two intersection points in red).	109
6.D.4	Deriving the deflection angle for a thin lens.	110
6.E.1	Deriving the geometric time delay.	113
7.2.1	Workflow of the targeted subthreshold search method (TESLA). . . .	121
7.3.1	The analysis flow of the simulation campaign.	122
7.3.2	Information about the mock data used for the simulation campaign. .	123

- 7.3.3 The templates in the original and targeted bank plotted in dark blue and orange respectively on the m_1 - m_2 plane. The best-match template for MGW220111a is indicated by a red star, and the true parameters of MGW220111a is represented by a green diamond. As we can see, even when the noise in the mock data is almost stationary and Gaussian, subthreshold lensed signals can still be found by templates with parameters very different from those within the posterior space of the superthreshold target event. This demonstrates that the posterior space of the target event itself is insufficient to cover all possible subthreshold lensed counterparts. 125
- 7.3.4 The templates in the original and “single-template” bank plotted in dark blue and orange respectively on the m_1 - m_2 plane. The “single-template” bank is a bank with 100 templates having component masses within $\pm 0.1 M_\odot$ from those of the posterior sample with maximum posterior to mimic the single template bank. 127
- 7.3.5 The templates in the original and PE bank, plotted in dark blue and orange respectively on the m_1 - m_2 plane. The best-match template for MGW220111a is indicated by a red star, and the true parameters of MGW220111a are represented by a green diamond. The purple curve represents the boundary to the 90% credible region of the posterior probability distribution for MGW220111a. The PE bank is generated by keeping only templates that lie within the 90% credible region of the posterior probability distribution, containing only 81 templates. . 128
- 7.3.6 The templates in the original and random bank plotted in dark blue and orange respectively on the m_1 - m_2 plane. The best-match template for MGW220111a is indicated by a red star, and the true parameters of MGW220111a are represented by a green diamond. The random bank contains the same number of templates, i.e. 552 templates, as the targeted template bank, and they are randomly selected from the original template bank. 129

- 7.3.7 (Top panel) The sensitive range v.s. FAR for MGW220111a-alike signals using the full template bank (black), targeted template bank (orange), “single-template” bank (blue), PE template bank (green) and random template bank (red) respectively. The shaded band for each curves represent the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. The percentage-change curve (orange) representing results using the targeted bank constructed by the TESLA search pipeline is above that of the full template bank, showing improvement in terms of sensitivity towards MGW220111a-alike (lensed) subthreshold signals. The vertical striped line shows the typical FAR threshold for triggers below which we consider as possible lensed candidates. 132
- 8.3.1 Merger rate density as a function of redshift based on the GWTC-2 results without lensing constraints (blue) and with lensing (red) included in LIGO–Virgo detections. We show the results for galaxy-scale lenses (G) and cluster-scale lenses (C) separately. Furthermore, S (or D) correspond to doubly lensed events where single (or double) events are detected. Because lensed detections occur at higher redshifts than unlensed events, their non-observation can be used to constrain mergers at higher redshifts. The results without lensing do not include constraints derived from the absence of a SGWB. 143
- 8.3.2 Cumulative fraction of lensed detectable mergers at any redshift with magnification greater than μ , constrained by the non-observation of the SGWB. The solid line shows the value obtained from the median merger rate density posterior. The shaded region corresponds to the 90% credible interval. Fewer than 1 in 10^3 events are expected to be lensed with magnification $\mu > 2$, on average. Significantly higher magnifications (e.g., $\mu > 30$) are suppressed by a further factor of 10. The results here show the probability of *observing* an event above a given magnification, which includes the merger-rate density and magnification bias information. 144

- 8.5.1 Scatter plot of the ranking statistics $\log_{10} \mathcal{B}^{\text{overlap}}$ and $\log_{10} \mathcal{R}^{\text{gal}}$ for a subset of event pairs that have both $\mathcal{B}^{\text{overlap}} > 50$ and $\mathcal{R}^{\text{gal}} > 0.01$. The dashed lines denote the significance levels of the combined ranking statistics (in terms of Gaussian standard deviations), obtained by simulating unlensed event pairs in Gaussian noise matching the O3a sensitivity of the LIGO–Virgo network. We identify several high $\mathcal{B}^{\text{overlap}} > 50$ candidates, which we follow up on with a detailed joint-PE analysis. We have used abbreviated event names, quoting the last 4 digits of the date identifier (see Table 8.5.1 for full names). 152
- 8.5.2 Combined results from the 39 sub-threshold searches with the GstLAL pipeline (top panel) and PyCBC pipeline (bottom panel). Each panel shows, as a solid line, the cumulative number of coincident triggers (observed) with inverse [s] greater than or equal to a given value. The dashed line is the expected distribution of background triggers, with the gray bands indicating uncertainties in multiples of the standard deviation σ of a Poisson distribution. For GstLAL, the results for this plot are obtained by a search over all O3a data using a combined bank from the 39 targeted banks. For PyCBC, the [s] are from the individual searches, but for triggers found by several of the single-template searches, their inverse [s] have been summed. . . . 164
- 8.6.1 The marginalized posterior distribution of redshifted lens mass M_L^z and $\log_{10} \mathcal{B}_U^{\text{ML}}$ between microlensed and unlensed hypotheses. The corresponding \log_{10} Bayes factors are noted to the right of the plot. We find no evidence of microlensing by point mass lenses. . . . 169
- 8.C.1 Distribution of microlensing Bayes factors $\log_{10} \mathcal{B}_U^{\text{ML}}$ for unlensed simulated signals, recovered using a lensed template. . . . 175
- 9.3.1 The False-Positive-Probabilities of each lensed candidate pair constructed from the set of gravitational wave events that exceed an astrophysical probability [159, 206] threshold of 0.5, as evaluated using the \mathcal{B}_U^{L} and Machine Learning classification statistics. Orange dashed lines that correspond to an False-Positive-Probability threshold of 10^{-2} , are also placed. Pairs whose \mathcal{B}_U^{L} -based or Machine Learning-based False-Positive-Probabilities fall below this threshold are selected for additional joint parameter estimation analyses. $\mathcal{B}_U^{\text{L}} < 10^{-6}$ has been mapped to an False-Positive-Probability of 1, which is reflected in the gap along the vertical axis between 0.4 and 1. 184

9.3.2	Bayes factors \mathcal{B}_U^L from <code>hanabi</code> for the highest-ranked multiple-image candidate pairs. As a check on the robustness of our results, we show the Bayes factors calculated using three different merger rate density models, namely the fiducial model tracking the Madau–Dickinson star-formation rate [239], and also the $R_{\min}(z)$ and $R_{\max}(z)$ model introduced in Abbott et al. [23]. The color for each marker represents the value of $p_{\text{astro}}^{\text{pair}}$ for each pair, which is the probability that both of the signals from a pair are of astrophysical origins and not from terrestrial sources.	187
9.3.3	Distribution of Bayes factors comparing different image type hypotheses for the 10 most relevant events. We compare the probability of being type II vs. type I (blue-solid histogram) and of being type II vs. type III (orange-dashed histogram). Only type II images display waveform distortions and for that reason, we do not compare type III vs. type I.	189
9.4.1	Distribution of microlensing \log_{10} Bayes factors $\mathcal{B}_U^{\text{Micro}}$ for all events in O3 (blue, solid line) and simulated unlensed signals (orange, dashed line) from Abbott et al. [23].	191
9.5.1	Merger rate density as a function of redshift based on the GWTC-3 results without lensing constraints (grey) and with lensing constraints (cross-hatching) included. For clarity, we show only the results for galaxy-scale lenses. Because lensed detections may occur at higher redshifts than unlensed events, their non-observation can be used to constrain the rate of mergers at higher redshifts. The ‘No lensing’ results shown here do not include constraints derived from the absence of an Stochastic Gravitational Wave Background. The latter constraints are shown separately by the solid black curves. . . .	195

- 9.5.2 The spread in the 90% upper limits on f_{DM} obtained from the O3 events using 5 different redshift distribution models for binary black holes mergers: Belczynski et al. [58], Dominik et al. [143], Madau and Dickinson [239], Abbott et al. [29] and uniform in comoving 4-volume, assuming a monochromatic mass spectrum for the compact objects forming dark matter. The lens mass is shown on the horizontal axis. The grey (black) shaded regions correspond to the spread in f_{DM} upper bounds computed assuming flat (Jeffreys) prior on Λ and Λ_ℓ . The upper and lower curves bounding the spreads correspond to the most pessimistic (weakest) and optimistic (strongest) upper limits, as determined from the set of assumed redshift distributions, in each mass bin. The current f_{DM} constraints are weaker relative to other corresponding EM constraints. We refer the reader to [91, 92, 112] for comparison. Nevertheless, these constraints are expected to improve significantly with the increased detection of unlensed BBHs in forthcoming LVK observing runs. 197
- 9.A.1 Posterior distribution of the Morse phase for GW200129_065458. We compare the real event posterior (solid-blue) with an injection campaign of type I (dashed-orange) and type II (dotted-green) images. Type II images correspond to $n_1 = 1/2$. For this event, the differences between the distribution are small and make it difficult to learn anything additional about the event. The Kolmogorov–Smirnov statistic is 0.07 for type I vs real, and 0.08 for type II vs real. 201
- 9.A.2 Posterior distribution of the Morse phase for GW190412. We compare the real event posterior (solid-blue) with an injection campaign of type I (dashed-orange) and type II (dotted-green) images. Type II images correspond to $n_1 = 1/2$. For this event, the peak seen in the real data and the one seen for the type II image are compatible, hinting at a possible type II image. In this case, the Kolmogorov–Smirnov statistic is 0.20 for type I vs real, and 0.13 for type II vs real. 202

9.A.3	Comparison of the Morse factor distribution for the real event (solid-purple) with the recovered posterior distribution for an injection made with IMRPhenomPv3HM for a type I image (dashed-blue) and a type II image (dotted-orange), and with an injection made with SEOBNRv4PHM for a type I image (dashed-green) and a type II image (dotted-red). In all the cases, the posterior distributions agree with the injected data, with the real event resembling a type II image.	203
9.B.1	Marginalized posterior distributions of redshifted lens mass M_L^z and $\log_{10} \mathcal{B}_U^{\text{Micro}}$ between microlensed and unlensed hypotheses.	204
9.B.2	The time-domain waveform corresponding to the maximum posterior of GW200208_130117, with and without the microlensing hypothesis.	205
10.3.1	Recovery of the Morse factor for the GW190412 event with different waveforms: IMRPhenomXPHM, IMRPhenomPv3HM, and SEOBNRv4PHM. The support for a type II image is present for the two waveforms from the Phenom family. However, the feature is less prominent for the IMRPhenomPv3HM waveform, and only marginally present for the SEOBNRv4PHM waveform. Therefore, the observed feature is probably spurious and the event is not a type II image.	213
10.3.2	Posteriors for a subset of the parameters including the detector frame chirp mass and mass ratio, the luminosity distance, the lens frame lens mass, and the dimensionless displacement of the source from the optical axis (i.e. the source position). These posteriors were produced during the point mass microlensing analysis done for GW190412 using GRAVELAMPS. As can be seen, similarly to GW191103 (shown in Fig. 10.4.5), the lensing parameter posteriors are extremely broad and uninformative. This is consistent with the expectations for a non-microlensing event.	215
10.4.1	Posteriors for some of the parameters obtained using the IMRPhenomXPHM waveform for GW191103 (blue) and GW191105 (orange). The overlap in the extrinsic parameters (e.g. sky location) is larger than that for the intrinsic parameters (e.g. detector-frame chirp mass and spins).	217

10.4.2	The highest ranked candidate strong lensing pairs from the PO analysis considering all the event pairs found based on the O3 data (dots) [32] and the supra-sub pair analyzed in this work. The dashed lines correspond to the 1σ and 2σ confidence levels for the combined PO statistic ($\mathcal{B}^{\text{overlap}} \times \mathcal{R}^{\text{gal}}$) with different lensing models computed from the background simulations. We note that beside GW191103–GW191105, the pair analyzed in this work, GW190728–GW190930 is also close to 1σ for PO. However, the pair has been discarded in previous searches with a lower overall significance than GW191103–GW191105 [23]. Therefore, it is not considered for further analyzes in this work.	218
10.4.3	Comparison of the posterior for the observed relative magnification and time delays for the O3 event pairs with the expected distributions for the lensed population of mergers from More and More [259] (dashed blue, using an SIE model) and from Wierda et al. [364] (solid green, using an elliptical model with a power law density profile and external shear (EPL-Shear)) and unlensed population (yellow-orange-red), both assuming galaxy lenses. Overlaid in brown are the observed values for selected O3 event pairs, and the letters mark the event pairs more compatible with the lensing hypothesis. Written in brown, and denoted with letter a, is a pair made of a super-threshold and a sub-threshold event, and further analysed in this work. The top panel corresponds to the expected distribution when the two images are of the same type, i.e., there is no phase difference between the two images (see top-left illustration), while the bottom panel corresponds to a configuration where the two image types differ, i.e., there is a $\pi/2$ shift between the images. Most of the observed event pairs are well outside the lensed distribution. The GW191103–GW191105, GW190706–GW190719, and GW101230–LGW200104 pairs are more compatible with the lensed hypothesis than with the unlensed one. In particular, the GW191103–GW191105 pair lies in a higher probability density region than the other pairs. One also sees that the GW191230–LGW200104 pair — made of a supra and a sub-threshold event — lies in a higher density region, even if it is less important than the GW191103–GW191105 pair. This pair is discussed further in Sec. 10.5.	245

10.4.4	From dark to lighter, the 10%, 50%, and 90% confidence sky localisation for the GW191103–GW191105 pair. Overlaid are the cross-matched 5 candidates from the Master Lens Database.	246
10.4.5	Posteriors for a subset of detector-frame source parameters and the lensing parameters produced during the GRAVELAMPS microlensing analysis of GW191103. As can be seen, whilst the source posteriors are well constrained, the lensing parameter posteriors are extremely broad and uninformative. This shows there are no observable microlensing features in this signal.	247
10.5.1	Scatter plot of posterior samples of GW191230, marginalised to the mass and aligned spin parameter space. We choose the template given by a KDE estimator. The other two options shown are the maximum likelihood point (ML) and an alternative KDE with an extra set of samples reflected on the $m_1 = m_2$ line (rKDE).	248
10.5.2	Overlaid LGW200104 and GW191230 skymaps with 90% and 50% confidence regions.	248
10.5.3	Posteriors for GW191230 (blue) and LGW200104 (orange). The posteriors, though broad, have significant overlap for both the intrinsic (left) and extrinsic (right) parameters.	249
10.5.4	The difference in Morse factor for the GW191230–LGW200104 event pair according to the GOLUM pipeline. The preferred value is 0.5 but there is also some support for 0.	250
10.5.5	Posterior distribution for the relative magnification for the GW191230–LGW200104 event pair measured with the GOLUM (solid blue) and HANABI (dashed orange) pipelines. One sees that the measured values are consistent between the two pipelines.	251
10.5.6	Posterior probability mass function for the image type of GW191230 and the image type of LGW200104 from HANABI. It is consistent with the GOLUM result that it is more likely for the difference in Morse factor to be $\Delta n = 0.5$ (i.e. the II-I and III-II configuration) than to be $\Delta n = 0$ (i.e. the I-I, II-II and III-III configuration).	252
10.5.7	Sky localisation 10%, 50%, and 90% confidence region (from dark to light) for the GW191230-LGW200104 pair. Overlaid are the cross-matched 21 candidates from the Master Lens Database.	252

10.6.1	Posteriors of a subset of source parameters as well as lensing parameters for GW200208 in the SIS microlensing case. Unlike the other events that have been examined within this work, the lensing parameters for the model are well constrained, even if the $3\text{-}\sigma$ confidence interval is a bit noisy. This means that this event, unlike the others, cannot be immediately ruled out as a lensing candidate from this test.	253
10.6.2	Posteriors of a subset of source parameters as well as lensing parameters for GW200208 in the point mass microlensing case. For this case, the posteriors are similarly constrained, notably arriving at similar lensing parameters as the SIS case even if the $3\text{-}\sigma$ confidence interval is a bit noisy.	254
10.6.3	The posterior densities of the recovered microlensing parameters for different parameter estimation runs, as presented in Table 10.6.1. The results are visualized with varying colors from light to dark (numbered from 1 to 11, indicating different runs as we move down the table.)	255
10.6.4	The residual power remaining, from top-to-bottom, in LIGO Hanford, LIGO Livingston, and Virgo when subtracting the best fit non-lensed waveform template for GW200208 as determined by the unlensed PE from the detector strain, over a subset of the total frequency range. As can be seen, there is no obviously coherent oscillatory behaviour in the residual power which would be expected in the case of the microlensing hypothesis. This absence would suggest that what remains is noise related rather than signal related.	255
10.6.5	Corner plot of the millilensing parameters obtained from a two-signal analysis of GW200208. Notable is that there is a peak within the effective luminosity distance and time delay parameters for a potential second millilensing signal. However, this peak could be explained by the presence of detector noise.	256
10.6.6	Discrete posterior of the number of signals in the multi-signal analysis for GW200208. This posterior is insufficient to confidently assert a number of signals present.	257

10.A.1	Comparison of sky maps for GW191105, as downloaded from GraceDB which is generated using LALINFERENCE (left) and from GWOSC which is generated using BILBY after de-glitching the Virgo data (right). The LALINFERENCE sky map is narrower as compared to BILBY, likely because of the glitch present in the data. Note that the localisation patch near to the north pole is the one that overlaps with GW191103 and therefore matters more than the rest of the patches.	259
10.A.2	Q-transforms images input to the LENSID pipeline for GW191103 (top panel) and GW191105 (bottom panel). The chirping feature for GW191105 is broken in both the LIGO detectors, whereas for GW191103 the chirp signal is fairly visible in Hanford, and not so visible in Livingston. This could be one of the reasons why LENSID did not identify this pair as significant.	260
10.B.1	Time-frequency map of the third-ranked PyCBC candidate. This shows the glitch present near to GPS time 1260223739. This represents the kind of quieter glitches that gets skipped in the normal autogating procedures.	260
10.B.2	Time-frequency map of the fourth-ranked PyCBC candidate, consistent with a scattered-light glitch [327, 328, 343].	261
10.C.1	Corner plot of the millilensing parameters obtained from a three-signal analysis. A similar peak appears in the second image parameters as was present within the two-signal analysis shown in Fig. 10.6.5. However, no such features are present within the posteriors of the third millilensing image.	262
10.C.2	Corner plots of the effective luminosity distance (left) and time delay (right) parameters obtained from the four-signal analysis. Consistent with the previous analyses, there are no peaking features in the third or fourth signal posteriors.	263
10.C.3	Corner plots of the effective luminosity distance (left) and relative time delays (right) obtained from a multi-signal analysis. Again, these are consistent with what has been seen in the previous analyses with no favouring for any number of signals above two.	264
11.3.1	Workflow of TESLA-X. Elements in grey are the same procedures as in traditional TESLA. Colored elements (color online) are new procedures in TESLA-X.	273

11.3.2	A cartoon depicting how we construct a reduced template bank from the injection results following the Traditional TESLA method (top right) and the proposed TESLA-X method (bottom left and right).	274
11.3.3	A cartoon depicting how we use the injection campaign results to construct a targeted population model in TESLA-X.	275
11.4.1	Rung-up templates by the simulated lensed injections of MS220508a in the injection campaign, plotted on the chirp mass \mathcal{M}_c - effective spin χ_{eff} space. The colors of the markers represent the network SNRs of the associated injections.	278
11.4.2	The templates in the full template bank (in grey) and in the TESLA reduced template bank (in orange), plotted in the \mathcal{M}_c - χ_{eff} space.	279
11.4.3	The contour map of the Gaussian Kernel Density Estimation function obtained for MS220508a in the \mathcal{M}_c - χ_{eff} space. The colors represent the probability density estimated by the KDE function.	280
11.4.4	The templates in the full template bank (in grey) and in the TESLA-X reduced template bank (in red), plotted in the \mathcal{M}_c - χ_{eff} space. The TESLA-X template bank is constructed by keeping templates that fall within the lowest contour level from the Gaussian KDE function shown in Figure 11.4.3, i.e. 0.005, plotted as a yellow curve in this figure. For easy comparison, we plot the TESLA reduced template bank templates in the same figure (in blue).	281
11.4.5	(Top panel) The sensitive range v.s. FAR for MS220508a-alike signals using the full template bank (black), TESLA template bank (red), and TESLA-X template bank (blue), respectively. The shaded band for each curve represents the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. We note that both the TESLA and TESLA-X bank improves in terms of sensitivity towards MS220508a-alike (lensed) signals, but the curve representing the TESLA-X bank (blue) is above that of the TESLA bank (red), meaning that the TESLA-X method outperforms the TESLA method. In particular, notice that the improvement is most significant in the region with $\text{FAR} \geq 10^{-7}$ Hz, which represents the “sub-threshold” region. This further proves that the TESLA-X method is better than the traditional TESLA method.	282

11.A.1	Rung-up templates by the simulated lensed injections of MS220508a in the injection campaign, plotted on the component masses m_1 - m_2 parameter space. The colors of the markers represent the network SNRs of the associated injections.	286
11.A.2	The templates in the full template bank (in grey) and in the TESLA reduced template bank (in orange), plotted in the component masses m_1 - m_2 space.	287
11.A.3	The contour map of the Gaussian Kernel Density Estimation function obtained for MS220508a in the m_1 - m_2 space. The colors represent the probability density estimated by the KDE function.	288
11.A.4	The templates in the full template bank (in grey) and in the “TESLA-X (component mass) bank” (in red), plotted in the m_1 - m_2 space. The “TESLA-X (component mass) bank” is constructed by keeping templates that fall within the lowest contour level from the Gaussian KDE function shown in Figure 11.A.3, i.e. 0.00002, plotted as a yellow curve in this figure. For easy comparison, we also plot the templates in the TESLA reduced template bank in the same figure (in blue).	289
11.A.5	(Top panel) The sensitive range v.s. FAR for MS220508a-alike signals using the full template bank (black), “TESLA-X (main) bank” (red), and “TESLA-X (component mass) bank” (blue) respectively. The shaded band for each curve represents the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. We note that while the “TESLA-X (main) bank” improves in terms of sensitivity towards MS220508a-alike (lensed) signals, “TESLA-X (component mass) bank” is performing worse than the full template bank, mainly due to the significant increase in trials factors caused by the large number of templates.	290
11.A.6	The contour map of the Gaussian Kernel Density Estimation function obtained for MS220510ae in the \mathcal{M}_c - χ_{eff} space. The colors represent the probability density estimated by the KDE function.	292
11.A.7	The contour map of the Gaussian Kernel Density Estimation function obtained for MS220510ae in the m_1 - m_2 space. The colors represent the probability density estimated by the KDE function.	293

11.A.8	(Top panel) The sensitive range v.s. FAR for MS220510ae-alike signals using the full template bank (black), “TESLA-X (main) bank” (red), and “TESLA-X (component mass) bank” (blue) respectively. The shaded band for each curve represents the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. We note that, unlike the case of MS220508a in the main text, the “TESLA-X (component mass) bank” improves in terms of sensitivity towards MS220510ae-alike (lensed) signals, “TESLA-X (main) bank” performs worse than the full template bank. This is mainly because spins are typically poorly measured for high chirp mass systems.	294
12.3.1	The Kernel Density Estimations of the distribution of log relative time delay and magnification ratios for the simulated lensed gravitational waves (green contours and pale blue sample points) for the SIE-double model and not-lensed, independent gravitational waves (black contours with grey sample points). We can see that the KDE for not-lensed, independent gravitational waves spans a larger range of relative magnifications and relative time delays solely because these gravitational-wave events are randomly distributed, in contrast to lensed gravitational waves under a given lens model.	300
12.4.1	The log relative time delay and magnification ratios of the rank 2 trigger (in red) and rank 3 trigger (in blue) respectively. The Kernel Density Estimations of the distribution of log relative time delay and magnification ratios for the simulated lensed gravitational waves (green contours) and not-lensed, independent gravitational waves (filled contours) are also plotted. We note that the rank 3 trigger, corresponding to the real sub-threshold lensed counterpart to MGW220111a, is slightly closer to the lensed KDE distribution than the rank 2 noise trigger, resulting in a slightly higher log lensing likelihood.	303

LIST OF TABLES

<i>Number</i>	<i>Page</i>
7.3.1	Information of the injected lensed gravitational-wave pair for the simulation campaign. All properties reported here are measured in the detector frame.
	123
7.3.2	Results of the targeted search of the simulation campaign for the two injected lensed signals.
	124
7.3.3	Results of the search for the two injected lensed signals using the random template bank.
	126
7.3.4	Results of the search for the two injected lensed signals using the PE template bank.
	127
7.3.5	Results of the search for the two injected lensed signals using the “single-template” bank.
	130
7.3.6	Number of injections found and missed during the search of mock data using the general template bank, TESLA targeted template bank, random template bank, PE template bank and “single-template” bank respectively.
	130
8.3.1	Expected fractional rates of observable lensed double events at current LIGO–Virgo sensitivity.
	141
8.4.1	Inferred properties of selected O3a events under the lensing magnification hypothesis.
	146
8.5.1	Summary of joint-PE results for event pairs in O3a.
	158
8.5.2	Candidates from individual sub-threshold searches for strongly-lensed counterpart images of the 39 O3a events from GWTC-2.
	165
9.3.1
	182
9.5.1	Expected fractional rates of observable lensed double events at current LIGO–Virgo sensitivity.
	193
10.4.1	PO statistic values for the GW191103–GW191105 pair using different waveform models in the single-event PE.
	220
10.4.2	Values of the detection statistic obtained using the GOLUM framework for the GW191103–GW191105 lensed candidate pair without lens model (C_U^L), and with an SIE lens, with ($C_{M_{\mu,t}}$) and without (C_{M_t}) relative magnification accounted for. The FAP_{PP} is decreased when using an SIE model.
	222

10.4.3	$\log_{10} \mathcal{B}_U^L$ for the GW191103–GW191105 pair from HANABI assuming three different merger rate density models and two different lens models. The values computed using the SIS model are reproduced from [32] for the sake of comparison. We see that the values with the SIE + external shear model are consistently higher than that with the SIS model, indicating a higher compatibility of the pair with a more realistic strong lensing model. However, since the values remain negative, the event is still most likely to be unlensed considering a more realistic lensing population with the most recent population models.	224
10.5.1	PyCBC targeted sub-threshold results for counterpart candidates to GW191230 ranked by IFAR. From left to right, the columns represent the event, the time delay compared to the supra-threshold event used to make the template, the inverse false-alarm rate (IFAR), the signal-to-noise ratio (SNR), and the 90% confidence region (CR) overlap for the sky posteriors.	229
10.5.2	Posterior-overlap factors for the GW191230–LGW200104 pair using different waveform models in the single-event PE.	230
10.5.3	Values of the detection statistic for the GW191230–LGW200104 lensed candidate pair without lens model (C_U^L), and with an SIE lens, with ($C_{M_{\mu,t}}$) and without (C_{M_t}) relative magnification accounted for. The FAP_{PP} is decreased when using an SIE model, showing that the observed characteristics are in line with the expected behaviour for the given model and population.	232
10.5.4	$\log_{10} \mathcal{B}_U^L$ for the GW191230 and LGW200104 pair from HANABI assuming three different (source) merger rate density models and two different lens models. We see that the values with the SIE + external shear model are all positive (but only mildly) and consistently higher than that with the SIS model which are all negative, indicating a higher compatibility of the pair with a more realistic strong lensing model. Note that the calculations assumed that both GW events are astrophysical of origin. These values are not sufficient to claim the event pair to be lensed as we would require a positive \log_{10} posterior odds, and the observed Bayes factors are not high enough to balance the low prior odds for strong lensing.	233

10.6.1	This table presents the results of a Bayesian model comparison study between the unlensed and the microlensed hypotheses for GW200208, with microlensing model corresponding to an isolated point-lens mass. The study was conducted for different configurations and sampler settings, as indicated by the different columns, to verify for possible noise artifacts and check the influence of the sampler settings on the results. The table includes the waveform approximant used, the lower and higher frequency cutoffs used for likelihood evaluation (f_{low} , f_{high}), duration of the data segment used, and the priors on the redshifted microlens mass (M_{Lz}^z) and the impact parameter (y), represented by $p(M_{Lz})$ and $p(y)$ respectively. The Bayes factor for the support of microlensing over the unlensed waveform model is given by $\log_{10} \mathcal{B}_U^L$. The range of the priors is also indicated. The terms ‘L.U’ and ‘L.L.U’ refer to log-normal and log-log-normal distributions respectively, while ‘P.L’ refers to a power law profile with the index given by α	236
10.6.2	Comparison of Bayes factors for the evidence against the unlensed hypothesis from the millilensing runs for GW200208.	241
11.4.1	Information of the injected super-threshold gravitational-wave signal MS220508a in the simulation campaign. All properties reported here are measured in the detector frame.	277
11.4.2	Number of injections found during the search of mock data using the general template bank, TESLA method, and TESLA-X method, respectively.	281
11.A.1	Number of injections found during the search of mock data using the general template bank, “TESLA-X (main) bank” and “TESLA-X (component mass) bank” respectively.	288
11.A.2	Information of the injected super-threshold gravitational-wave signal MS220510ae in the simulation campaign. All properties reported here are measured in the detector frame.	291
11.A.3	Number of injections found during the search of mock data using the general template bank, “TESLA-X (main) bank” and “TESLA-X (component mass) bank” respectively for the mock event MS220510ae.	292
12.4.1	Top three candidates from the TESLA targeted search for the mock lensed event MGW220111a.	301

12.4.2	Log Lensing Likelihoods for the rank 2 and 3 triggers evaluated based on the SIE-double lens model.	302
--------	--	-----

Part I

Introduction to gravitational waves and gravitational-wave data analysis with the GstLAL Search pipeline

Chapter 1

OVERVIEW OF THE THESIS

This thesis is divided into two parts: Part I - “Introduction to gravitational waves and gravitational-wave data analysis with the GstLAL Search pipeline” and Part II - “Gravitational Lensing of Gravitational Waves”.

Part I is structured as follows: In Chapter 2, I give an overview of general relativity, gravitational waves, and the Advanced LIGO detectors. In Chapter 3, I go over briefly the major steps, namely (1) data calibration, (2) detector characterization and noise mitigation, (3) search and (4) Bayesian parameter estimation, in the LIGO-Virgo-KAGRA’s data analysis framework to search for gravitational waves, specifically from compact binary coalescences. In Chapter 4, I first give a top-level introduction to the general working principle of modelled matched-filtering-based search pipelines for gravitational waves from compact binary coalescences. I then focus on the GstLAL search pipeline, which will be the primary search pipeline to be used and discussed in various scenarios in the rest of the thesis. Alongside reviewing the steps of a typical GstLAL analysis, I selected a few key terms in the likelihood ratio calculation from the GstLAL search pipeline for discussion. Some of these terms will be tuned/modified in later chapters for specific purposes. In Chapter 5, I propose a method to improve the search sensitivity of the GstLAL search pipeline by properly using statistical data quality information generated by a machine-learning-based pipeline iDQ. Since the work is ongoing, I will only explain and outline the expected results should the method be correctly implemented into GstLAL.

In Part II, we focus on a specific phenomenon, “gravitational lensing”, as predicted by general relativity. In Chapter 6, I give a detailed introduction to gravitational lensing and lensing of gravitational waves. I also outline the current effort by the LVK collaboration to search for lensing signatures in gravitational waves from the third observing run O3. Chapter 7 is an adaption of [223], where we introduce a GstLAL-based targeted search method (TESLA) to look for possible strongly-lensed sub-threshold counterparts to known superthreshold gravitational-wave events. Chapter 8, 9 and 10 present results from the LVK-collaboration-wide effort to search for gravitationally-lensed gravitational waves from the third ob-

serving run O3. The search for sub-threshold lensed gravitational waves using the TESLA pipeline discussed in Chapter 7 is also included in the analyses. In Chapter 11, we explain how we upgrade the traditional TESLA pipeline introduced in Chapter 7 to the TESLA-X pipeline with improved search sensitivity towards strongly-lensed sub-threshold gravitational waves. This is done by a better approach to constructing targeted template banks and using targeted population models. Chapter 12 discusses yet another upgrade made to the TESLA-X pipeline. In the chapter, we explain how we assign an alternative ranking statistic, the Lensing Likelihood $\mathcal{L}_{\text{lensing}}$, to each trigger reported by the TESLA-X pipeline. The Lensing Likelihood is evaluated based on the consistency of the triggers to an assumed lens model. This alternative ranking statistic serves as additional information to determine which candidates are more likely to be sub-threshold strongly-lensed counterparts to their targets and their priority for running subsequent lensing follow-up analyses. The upgraded TESLA-X pipeline will continue to be one of the key pipelines of the LVK collaboration to search for possible strongly-lensed gravitational waves in the fourth observing run O4.

Finally, in Chapter 13, we summarize the results and findings in this thesis and outline the possible future work.

Chapter 2

INTRODUCTION TO GRAVITATIONAL WAVES

2.1 Overview of this Chapter

In this introductory chapter, I will set the stage for our discussion in later chapters by briefly reviewing the theory of general relativity, gravitational waves, and the Advanced LIGO detectors. This chapter is structured as follows: I will first give a top-level overview of general relativity in Section 2.2. In Section 2.3, I will give the standard discussion and derivations of linearized gravity in the weak-field limit. Then, in Section 2.4, I explain how gravitational waves emerge from the resulting Einstein's field equations in the weak-field limit. Section 2.5 derives the energy carried by gravitational waves. Finally, Section 2.6 introduces the Advanced LIGO detectors and derives the antenna pattern functions associated with the L-shape geometry of the LIGO detectors. The antenna pattern functions are essential for evaluating the input to the LIGO detectors as a gravitational wave passes through.

2.2 A very brief overview of General Relativity

The most important equation in General Relativity is the Einstein's Field equation, given by

$$G^{\mu\nu} = 8\pi G T^{\mu\nu}. \quad (2.1)$$

Here, the quantities with superscripts are tensors. Each index (i.e. μ and ν) can take 4 values, representing time (1 dimension) and space (3 dimensions). G is the gravitational constant.

Let us briefly go over each quantity. $G^{\mu\nu}$ is known as the Einstein tensor, defined as

$$G^{\mu\nu} = R^{\mu\nu} - \frac{1}{2} g^{\mu\nu} R, \quad (2.2)$$

where $R^{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar, and $g^{\mu\nu}$ is the metric tensor.

Now, the Ricci scalar is simply a contraction of the Ricci tensor:

$$R = g^{\mu\nu} R_{\mu\nu}. \quad (2.3)$$

The Ricci tensor is a contraction of yet another tensor called the Riemann tensor:

$$R_{\alpha\beta} = R^{\kappa}{}_{\alpha\kappa\beta} = g^{\mu\kappa} R_{\mu\alpha\kappa\beta}. \quad (2.4)$$

The Riemann tensor is defined as

$$R^\alpha_{\beta\mu\nu} = \partial_\mu \Gamma^\alpha_{\beta\nu} - \partial_\nu \Gamma^\alpha_{\beta\mu} + \Gamma^\alpha_{\mu\gamma} \Gamma^\gamma_{\beta\nu} - \Gamma^\alpha_{\nu\gamma} \Gamma^\gamma_{\beta\mu}, \quad (2.5)$$

where the Γ 's are the Christoffel symbols, defined as

$$\Gamma^\sigma_{\alpha\beta} = \frac{1}{2} g^{\sigma\rho} \left[\partial_\alpha g_{\rho\beta} + \partial_\beta g_{\rho\alpha} - \partial_\rho g_{\alpha\beta} \right]. \quad (2.6)$$

We now arrive at an expression that depends only on the metric tensor, i.e. the geometry of the spacetime we are considering.

That being said, the Einstein tensor, $G^{\mu\nu}$, is solely a geometrical quantity encoding the geometry of spacetime we are considering.

On the other hand, the tensor $T^{\mu\nu}$, known as the Stress-Energy tensor, is defined as

$$T^{\mu\nu} = (\rho_0 + P_0) u^\nu u^\mu + P_0 g^{\mu\nu}, \quad (2.7)$$

where ρ_0 is the density, P_0 is the pressure, and u^μ is the four-velocity. In short, the "stress-energy tensor" is related to the energy of and forces on the object under consideration.

Essentially, Einstein's field equations tell us that the (gravitational) force on an object and its energy are closely related to the geometry of the spacetime in which it is situated. John Wheeler gave a precise and concise statement for this equation:

"Matter tells spacetime how to curve, and spacetime tells matter how to move."

2.3 Linearized gravity in the weak-field limit

In flat Minkowski's spacetime, we know that the metric tensor is given by the Minkowski's metric as

$$\eta_{\mu\nu} = \text{diag}[-1, 1, 1, 1]. \quad (2.8)$$

We now consider the weak-field limit ¹ in which the metric tensor, encoding the geometry in the region, is only the Minkowski metric $\eta_{\mu\nu}$ plus a small additional symmetric perturbation term $h_{\mu\nu}$, i.e.

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}. \quad (2.9)$$

¹ By weak-field, we want to consider a region of spacetime that is only slightly curved (e.g. this could be a region far away from a black hole).

Note that when we say the perturbation term is small, it means that

$$|h_{\mu\nu}| \ll 1. \quad (2.10)$$

We can now evaluate the Christoffel symbol, defined as

$$\Gamma^\sigma_{\alpha\beta} = \frac{1}{2}g^{\sigma\rho}(\partial_\alpha g_{\rho\beta} + \partial_\beta g_{\rho\alpha} - \partial_\rho g_{\alpha\beta}) \quad (2.11)$$

$$= \frac{1}{2}g^{\sigma\rho}(g_{\rho\beta,\alpha} + g_{\rho\alpha,\beta} - g_{\alpha\beta,\rho}), \quad (2.12)$$

where in the last line, I have replaced the “partial derivative” expression $\partial_\alpha K$ with a simplified expression $K_{,\alpha}$, with the perturbed metric tensor.

The parentheses in the Christoffel symbol expression include terms of partial derivatives of the metric tensors, e.g. $g_{\alpha\beta,\rho}$. Since the Minkowski metric is constant, each of these expressions reduces to the derivative of the perturbation term only, i.e.

$$g_{\alpha\beta,\rho} = \underbrace{\eta_{\alpha\beta,\rho}}_{=0} + h_{\alpha\beta,\rho} = h_{\alpha\beta,\rho}. \quad (2.13)$$

We then have

$$\Gamma^\sigma_{\alpha\beta} = \frac{1}{2}g^{\sigma\rho}(g_{\rho\beta,\alpha} + g_{\rho\alpha,\beta} - g_{\alpha\beta,\rho}) \quad (2.14)$$

$$= \frac{1}{2}(\eta^{\sigma\rho} - h^{\sigma\rho})(h_{\rho\beta,\alpha} + h_{\rho\alpha,\beta} - h_{\alpha\beta,\rho}). \quad (2.15)$$

As $h_{\alpha\beta}$ is small, we neglect terms that are of order higher than $h_{\alpha\beta}$. This gives

$$\Gamma^\sigma_{\alpha\beta} = \frac{1}{2}(\eta^{\sigma\rho} - h^{\sigma\rho})(h_{\rho\beta,\alpha} + h_{\rho\alpha,\beta} - h_{\alpha\beta,\rho}) \quad (2.16)$$

$$\approx \frac{1}{2}\eta^{\sigma\rho}(h_{\rho\beta,\alpha} + h_{\rho\alpha,\beta} - h_{\alpha\beta,\rho}). \quad (2.17)$$

Now, the Riemann Tensor is defined as

$$R^\alpha_{\beta\mu\nu} = \Gamma^\alpha_{\beta\nu,\mu} - \Gamma^\alpha_{\beta\mu,\nu} + \Gamma^\alpha_{\mu\gamma}\Gamma^\gamma_{\beta\nu} - \Gamma^\alpha_{\nu\gamma}\Gamma^\gamma_{\beta\mu}. \quad (2.18)$$

We notice that the last two terms in the original Riemann Tensor expression involve multiplication of two Christoffel symbols. Previously we have found that the Christoffel symbols are of the same order as $h_{\mu\nu}$, and so when two of them are

multiplied together, they become “higher order terms” which we can neglect. Hence we have

$$R^\alpha_{\beta\mu\nu} \approx \Gamma^\alpha_{\beta\nu,\mu} - \Gamma^\alpha_{\beta\mu,\nu} \quad (2.19)$$

$$= \frac{1}{2}\eta^{\alpha\rho} \left[\left(h_{\rho\nu,\beta} + h_{\rho\beta,\nu} - h_{\beta\nu,\mu} \right)_{,\mu} - \left(h_{\rho\mu,\beta} + h_{\rho\beta,\mu} - h_{\beta\mu,\nu} \right)_{,\nu} \right] \quad (2.20)$$

$$= \frac{1}{2}\eta^{\alpha\rho} \left[h_{\rho\nu,\beta\mu} + h_{\beta\mu,\rho\nu} - h_{\rho\mu,\beta\nu} - h_{\beta\mu,\rho\nu} + \underbrace{(h_{\rho\beta,\mu\nu} - h_{\rho\beta,\mu\nu})}_{=0} \right] \quad (2.21)$$

$$= \frac{1}{2}\eta^{\alpha\rho} \left[h_{\rho\nu,\beta\mu} + h_{\beta\mu,\rho\nu} - h_{\rho\mu,\beta\nu} - h_{\beta\mu,\rho\nu} \right]. \quad (2.22)$$

The Ricci Tensor is defined as

$$R_{\alpha\beta} = R^\mu_{\alpha\mu\beta} = g^{\mu\nu} R_{\nu\alpha\mu\beta}. \quad (2.23)$$

For the perturbed metric, we get

$$R_{\alpha\beta} = g^{\mu\nu} R_{\nu\alpha\mu\beta} \quad (2.24)$$

$$= \frac{1}{2}g^{\mu\nu} \left[h_{\nu\beta,\alpha\mu} + h_{\alpha\mu,\nu\beta} - h_{\nu\mu,\alpha\beta} - h_{\alpha\beta,\nu\mu} \right] \quad (2.25)$$

$$\approx \frac{1}{2}\eta^{\mu\nu} \left[h_{\nu\beta,\alpha\mu} + h_{\alpha\mu,\nu\beta} - h_{\nu\mu,\alpha\beta} - h_{\alpha\beta,\nu\mu} \right] \quad (2.26)$$

$$= \frac{1}{2} \left[h^\mu_{\beta,\alpha\mu} + h_{\alpha\mu}{}^\mu{}_\beta - h_{,\alpha\beta} - \square^2 h_{\alpha\beta} \right], \quad (2.27)$$

where we have simplified the expressions with

$$K = K^\mu_\mu = g^{\mu\nu} K_{\mu\nu} \quad (2.28)$$

that represents the trace of the tensor $K_{\mu\nu}$, and

$$\square^2 = \partial^\mu \partial_\mu = -\partial^2 t + \partial^2 x + \partial^2 y + \partial^2 z \quad (2.29)$$

that represents the D'Alembertian operator.

The Ricci scalar, as mentioned before, is simply the contraction of the Ricci Tensor, defined by

$$R = g^{\mu\nu} R_{\mu\nu}, \quad (2.30)$$

which, together with the Ricci Tensor expression, gives

$$R = g^{\alpha\beta} R_{\alpha\beta} = \frac{1}{2} g^{\alpha\beta} \left[h^\mu_{\beta,\alpha\mu} + h_{\alpha\mu}{}^{,\mu}{}_\beta - h_{,\alpha\beta} - \square^2 h_{\alpha\beta} \right] \quad (2.31)$$

$$\approx \frac{1}{2} \eta^{\alpha\beta} \left[h^\mu_{\beta,\alpha\mu} + h_{\alpha\mu}{}^{,\mu}{}_\beta - h_{,\alpha\beta} - \square^2 h_{\alpha\beta} \right] \quad (2.32)$$

$$= \frac{1}{2} \left[h^{\mu\alpha}{}_{,\alpha\mu} + h_{\alpha\mu}{}^{,\mu\alpha} - h_{,\alpha}{}^\alpha - \square^2 h_\alpha{}^\alpha \right] \quad (2.33)$$

$$= \frac{1}{2} \left[2h^{\mu\alpha}{}_{,\mu\alpha} - 2\square^2 h \right] \quad (2.34)$$

$$= h^{\mu\alpha}{}_{,\mu\alpha} - \square^2 h, \quad (2.35)$$

where we have made use of the facts that (1) $K_{\alpha\beta}{}^{,\alpha\beta} = K^{\alpha\beta}{}_{,\alpha\beta}$, (2) that partial derivatives are commutable, and (3) that the perturbation term is symmetric.

We are now ready to evaluate the Einstein Tensor, defined as

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R. \quad (2.36)$$

which, with the expressions above, can be expressed as

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R \quad (2.37)$$

$$= \frac{1}{2} \left[h^\alpha_{\nu,\mu\alpha} + h^\alpha_{\mu,\nu\alpha} - h_{,\mu\nu} - \square^2 h_{\mu\nu} \right] - \frac{1}{2} g_{\mu\nu} \left[h^{\alpha\beta}{}_{,\alpha\beta} - \square^2 h \right] \quad (2.38)$$

$$= \frac{1}{2} \left[h^\alpha_{\nu,\mu\alpha} + h^\alpha_{\mu,\nu\alpha} - h_{,\mu\nu} - \square^2 h_{\mu\nu} - \eta_{\mu\nu} h^{\alpha\beta}{}_{,\alpha\beta} + \eta_{\mu\nu} \square^2 h \right]. \quad (2.39)$$

By considering the term

$$H_{\mu\nu} = h_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} h, \quad (2.40)$$

whose trace is

$$H = \eta^{\mu\nu} H_{\mu\nu} \quad (2.41)$$

$$= \eta^{\mu\nu} h_{\mu\nu} - \frac{1}{2} \eta^{\mu\nu} \eta_{\mu\nu} h \quad (2.42)$$

$$= h - \frac{1}{2} (4) h \quad (2.43)$$

$$= -h, \quad (2.44)$$

we can write

$$H_{\mu\nu} = h_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} h = h_{\mu\nu} + \frac{1}{2} \eta_{\mu\nu} H. \quad (2.45)$$

By rearranging the terms, we obtain

$$h_{\mu\nu} = H_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}H. \quad (2.46)$$

We can then rewrite the Einstein tensor as

$$G_{\mu\nu} = \frac{1}{2} \left[h^\alpha_{\nu,\mu\alpha} + h^\alpha_{\mu,\nu\alpha} - h_{,\mu\nu} - \square^2 h_{\mu\nu} - \eta_{\mu\nu} h^{\alpha\beta}_{,\alpha\beta} + \eta_{\mu\nu} \square^2 h \right] \quad (2.47)$$

$$= \frac{1}{2} \left[\left(H^\alpha_{\nu,\mu\alpha} - \frac{1}{2}\eta^\alpha_\nu H_{,\mu\alpha} \right) + \left(H^\alpha_{\mu,\nu\alpha} - \frac{1}{2}\eta^\alpha_\mu H_{,\nu\alpha} \right) + H_{,\mu\nu} \right] \quad (2.48)$$

$$+ \frac{1}{2} \left[- \left(\square^2 H_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu} \square^2 H \right) - \eta_{\mu\nu} \left(H^{\alpha\beta}_{,\alpha\beta} - \frac{1}{2}\eta^{\alpha\beta} H_{,\alpha\beta} \right) - \eta_{\mu\nu} \square^2 H \right] \quad (2.49)$$

$$= \frac{1}{2} \left[H^\alpha_{\nu,\mu\alpha} - \frac{1}{2}H_{,\mu\nu} + H^\alpha_{\mu,\nu\alpha} - \underbrace{\frac{1}{2}H_{,\nu\mu}}_{=\frac{1}{2}H_{,\mu\nu}} + H_{,\mu\nu} \right] \quad (2.50)$$

$$+ \frac{1}{2} \left[-\square^2 H_{\mu\nu} + \frac{1}{2}\eta_{\mu\nu} \square^2 H - \eta_{\mu\nu} H^{\alpha\beta}_{,\alpha\beta} + \frac{1}{2}\eta_{\mu\nu} \underbrace{H_{,\alpha}{}^\alpha}_{=\square^2 H} - \eta_{\mu\nu} \square^2 H \right] \quad (2.51)$$

$$= \frac{1}{2} \left[H^\alpha_{\nu,\mu\alpha} + H^\alpha_{\mu,\nu\alpha} - \square^2 H_{\mu\nu} - \eta_{\mu\nu} H^{\alpha\beta}_{,\alpha\beta} \right] \quad (2.52)$$

$$= \frac{1}{2} \left[H_{\alpha\nu,\mu}{}^\alpha + H_{\alpha\mu,\nu}{}^\alpha - \square^2 H_{\mu\nu} - \eta_{\mu\nu} H_{\alpha\beta,}{}^{\alpha\beta} \right]. \quad (2.53)$$

With gauge freedom, we can freely transform potentials via coordinate transformation without altering the final results. The idea is that we shift the coordinates x^α to a new coordinate \tilde{x}^α , each by a small amount ξ^α , i.e.

$$\tilde{x}^\alpha = x^\alpha + \xi^\alpha, \quad (2.54)$$

where $\xi^\alpha = \xi^\alpha(x^\beta)$ is a function of x^β assuming that $|\xi^\alpha| \ll 1$. Note that

$$\frac{\partial x^\alpha}{\partial \tilde{x}^\sigma} = \frac{\partial \tilde{x}^\alpha}{\partial \tilde{x}^\sigma} - \frac{\partial \xi^\alpha}{\partial \tilde{x}^\sigma} \quad (2.55)$$

$$= \delta^\alpha_\sigma - \xi^\alpha_{,\beta} \frac{\partial x^\beta}{\partial \tilde{x}^\sigma} \quad (2.56)$$

$$= \delta^\alpha_\sigma - \xi^\alpha_{,\beta} \left(\delta^\beta_\sigma - \frac{\partial \xi^\beta}{\partial \tilde{x}^\sigma} \right) \quad (2.57)$$

$$= \delta^\alpha_\sigma - \xi^\alpha_{,\sigma} + \underbrace{\xi^\alpha_{,\beta} \frac{\partial \xi^\beta}{\partial \tilde{x}^\sigma}}_{H.O.T.} \quad (2.58)$$

$$\approx \delta^\alpha_\sigma - \xi^\alpha_{,\sigma} \quad (2.59)$$

and

$$\frac{\partial \tilde{x}^\alpha}{\partial x^\sigma} = \frac{\partial x^\alpha}{\partial x^\sigma} + \frac{\partial \xi^\alpha}{\partial \tilde{x}^\sigma} \quad (2.60)$$

$$= \delta^\alpha_\sigma + \xi^\alpha_{,\beta} \frac{\partial \tilde{x}^\beta}{\partial x^\sigma} \quad (2.61)$$

$$= \delta^\alpha_\sigma + \xi^\alpha_{,\beta} \left[\delta^\beta_\sigma + \frac{\xi^\beta}{\partial x^\sigma} \right] \quad (2.62)$$

$$\approx \delta^\alpha_\sigma + \xi^\alpha_{,\sigma} \quad (2.63)$$

where we have neglected higher-order terms (listed as ‘‘H.O.T.’’ in the equations).

With the coordinate transformation, the metric tensor transforms as

$$\tilde{g}_{\mu\nu} = \Lambda_{\mu'}^\alpha \Lambda_{\nu'}^\beta g_{\alpha\beta} \quad (2.64)$$

$$= \frac{dx^\alpha}{d\tilde{x}^\mu} \frac{dx^\beta}{d\tilde{x}^\nu} g_{\alpha\beta} \quad (2.65)$$

$$\approx [\delta^\alpha_\mu - \xi^\alpha_{,\mu}] [\delta^\beta_\nu - \xi^\beta_{,\nu}] g_{\alpha\beta} \quad (2.66)$$

$$= g_{\mu\nu} - g_{\mu\beta} \xi^\beta_{,\nu} - g_{\alpha\nu} \xi^\alpha_{,\mu} + \text{H.O.T.} \quad (2.67)$$

$$\approx g_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu}, \quad (2.68)$$

and because Minkowski’s metric $\eta_{\mu\nu}$ is a constant tensor regardless of what coordinate system we are using, we have

$$\tilde{g}_{\mu\nu} + \eta_{\mu\nu} = g_{\mu\nu} + \eta_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu} \quad (2.69)$$

and hence the perturbation term $h_{\mu\nu}$ becomes

$$\tilde{h}_{\mu\nu} = h_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu}. \quad (2.70)$$

Note that

$$\tilde{h} = \eta^{\alpha\beta} \tilde{h}_{\alpha\beta} \quad (2.71)$$

$$= \eta^{\alpha\beta} \left[h_{\alpha\beta} - \xi_{\alpha,\beta} - \xi_{\beta,\alpha} \right] \quad (2.72)$$

$$= h - \xi^{\beta}_{,\beta} - \underbrace{\xi^{\alpha}_{,\alpha}}_{=\xi^{\beta}_{,\beta}} \quad (2.73)$$

$$= h - 2\Box \cdot \vec{\xi}. \quad (2.74)$$

Hence we get

$$\tilde{H}_{\mu\nu} = \tilde{h}_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\tilde{h} \quad (2.75)$$

$$= [h_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu}] - \frac{1}{2}\eta_{\mu\nu} [h - 2\Box \cdot \vec{\xi}] \quad (2.76)$$

$$= \underbrace{\left(h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h \right)}_{=H_{\mu\nu}} - \xi_{\mu,\nu} - \xi_{\nu,\mu} + \eta_{\mu\nu}\Box \cdot \vec{\xi} \quad (2.77)$$

$$= H_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu} + \eta_{\mu\nu}\Box \cdot \vec{\xi}. \quad (2.78)$$

We notice that

$$\tilde{R}_{\alpha\beta\mu\nu} = \frac{1}{2} \left[\tilde{h}_{\alpha\nu,\beta\mu} + \tilde{h}_{\beta\mu,\alpha\nu} - \tilde{h}_{\alpha\mu,\beta\nu} - \tilde{h}_{\beta\nu,\alpha\mu} \right] \quad (2.79)$$

$$= R_{\alpha\beta\mu\nu} + \frac{1}{2} \left[-\xi_{\alpha,\nu\beta\mu} - \xi_{\nu,\alpha\beta\mu} - \xi_{\beta,\mu\alpha\nu} - \xi_{\mu,\beta\alpha\nu} + \xi_{\alpha,\mu\beta\nu} + \xi_{\mu,\alpha\beta\nu} + \xi_{\beta,\nu\alpha\mu} + \xi_{\nu,\beta\alpha\mu} \right] \quad (2.80)$$

$$= R_{\alpha\beta\mu\nu}. \quad (2.81)$$

This means that the Riemann tensor is not affected by gauge transformation up to the first order, and hence, we are free to choose a “friendly” coordinate system to simplify our manipulation of the Einstein’s Field equations without affecting the physical meaning of the solutions.

In terms of $H_{\mu\nu}$ and H , the Einstein’s Field equations read

$$\frac{1}{2} \left[H_{\alpha\nu,\mu}{}^{\alpha} + H_{\alpha\mu,\nu}{}^{\alpha} - \Box^2 H_{\mu\nu} - \eta_{\mu\nu} H_{\alpha\beta,}{}^{\alpha\beta} \right] = 8\pi G T_{\mu\nu}. \quad (2.82)$$

Note that three out of four of the terms on the left-hand side are in the form

$$H_{\mu\nu},{}^{\mu}. \quad (2.83)$$

(Note that the terms involve a second partial derivative, but here, we are just focusing on the partial derivative with an index matching one of that of the tensors being differentiated.)

By gauge transformation, with the transformed perturbation term $\tilde{H}_{\mu\nu}$, we have

$$\tilde{H}_{\mu\nu} = H_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu} + \eta_{\mu\nu} \square \cdot \vec{\xi} \quad (2.84)$$

$$\tilde{H}_{\mu\nu}{}^{;\mu} = 0 = H_{\mu\nu}{}^{;\mu} - \underbrace{\xi_{\mu,\nu}{}^{;\mu}}_{=\partial_\nu(\partial^\mu \xi_\mu)=0} - \underbrace{\xi_{\nu,\mu}{}^{;\mu}}_{=\square^2 \xi_\nu} + \underbrace{\partial^\mu(\eta_{\mu\nu} \square \cdot \vec{\xi})}_{=0} \quad (2.85)$$

$$0 = H_{\mu\nu}{}^{;\mu} - \square^2 \xi_\nu \quad (2.86)$$

$$\square^2 \xi_\nu = H_{\mu\nu}{}^{;\mu}. \quad (2.87)$$

This means we can always find a gauge ξ^α satisfying

$$\square^2 \xi^\alpha = H^{\alpha\beta}{}_{;\beta}, \quad (2.88)$$

such that we can transform $H_{\mu\nu}$ into $\tilde{H}_{\mu\nu}$ that gives

$$\tilde{H}_{\mu\nu}{}^{;\mu} = 0. \quad (2.89)$$

With this gauge transformation (formally known as the Lorenz gauge), we can write Einstein's field equations as

$$\frac{1}{2} [-\square^2 H_{\mu\nu}] = 8\pi G T_{\mu\nu} \quad (2.90)$$

$$\square^2 H_{\mu\nu} = -16\pi G T_{\mu\nu}. \quad (2.91)$$

We have not yet exhausted all the gauge freedom. Suppose now we are in a coordinate system such that

$$H_{\mu\nu}{}^{;\mu} = 0 \quad (2.92)$$

already. Now consider we do yet another coordinate transformation

$$\tilde{x}^\alpha = x^\alpha + \xi^\alpha, \quad (2.93)$$

such that $\tilde{h}_{\mu\nu}$ and $\tilde{H}_{\mu\nu}$ are transformed accordingly. Note that we still have

$$\tilde{H}_{\mu\nu} = H_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu} + \eta_{\mu\nu} \square \cdot \vec{\xi}. \quad (2.94)$$

If we now apply the partial derivative ∂^μ on both sides, we get

$$\tilde{H}_{\mu\nu}{}^{;\mu} = 0 = \underbrace{H_{\mu\nu}{}^{;\mu}}_{=0} - \underbrace{\xi_{\mu,\nu}{}^{;\mu}}_{=\partial_\nu \square \cdot \vec{\xi}=0} - \xi_{\nu,\mu}{}^{;\mu} + \underbrace{\partial^\mu(\eta_{\mu\nu} \square \cdot \vec{\xi})}_{=0} \quad (2.95)$$

$$\square^2 \xi_\nu = 0. \quad (2.96)$$

Hence, as long as the gauge ξ_ν satisfies the above condition, Einstein's field equations will remain invariant.

2.4 Gravitational waves

The Einstein's field equations, after applying the Lorenz gauge, is given by

$$\square^2 H_{\mu\nu} = -16\pi G T_{\mu\nu}. \quad (2.97)$$

In empty space, there is no mass/energy, and hence the stress-energy tensor on the right becomes zero, which gives

$$\square^2 H_{\mu\nu} = 0. \quad (2.98)$$

Recall that the \square^2 actually means

$$\square^2 = -\partial_t^2 + \nabla^2 = -\partial_{tt} + \nabla^2, \quad (2.99)$$

and, assuming we can treat $H_{\mu\nu}$ simply as a function, the Einstein's field equation reads

$$\partial_{tt} H_{\mu\nu} = \nabla^2 H_{\mu\nu} = (\partial_{xx} + \partial_{yy} + \partial_{zz}) H_{\mu\nu}, \quad (2.100)$$

where in the last line I assumed we use Cartesian coordinates for spatial coordinates. This equation resembles the wave equation we see in classical mechanics. In the same spirit, we can expect plane wave-like solutions to the empty space Einstein's field equations. We assume the solution to be in the form

$$H_{\mu\nu} = A_{\mu\nu} \cos(-\omega t + \vec{k} \cdot \vec{r}) = A_{\mu\nu} \cos(k_\alpha x^\alpha) = A_{\mu\nu} \cos(g_{\alpha\beta} k^\alpha x^\beta), \quad (2.101)$$

where $k_\alpha = [-\omega, k_x, k_y, k_z]$ is the four wave co-vector, and $A_{\mu\nu}$ is just a constant matrix (tensor).

Several constraints exist on the constant tensor $A_{\mu\nu}$ and the four wave-co-vector k^α .

(1) Recall that $H_{\mu\nu}$, which is a function of $h_{\mu\nu}$, is symmetric (because the metric $g_{\mu\nu}$ is symmetric), i.e.

$$H_{\mu\nu} = H_{\nu\mu}. \quad (2.102)$$

Therefore, we must have

$$A_{\mu\nu} = A_{\nu\mu}. \quad (2.103)$$

That is, $A_{\mu\nu}$ must also be symmetric. (2) Also, $H_{\mu\nu}$ needs to satisfy the Lorenz gauge, i.e.

$$H_{\mu\nu}{}^{;\mu} = 0. \quad (2.104)$$

We hence have

$$H_{\mu\nu}{}^{;\mu} = \frac{\partial}{\partial x^\mu} [A_{\mu\nu} \cos(k_\alpha x^\alpha)] = 0. \quad (2.105)$$

$$A_{\mu\nu} \frac{\partial}{\partial x^\mu} [\cos(k_\alpha x^\alpha)] = 0 \quad (2.106)$$

$$-A_{\mu\nu} \sin(k_\beta x^\beta) k_\alpha \underbrace{\frac{\partial x^\alpha}{\partial x^\mu}}_{\delta^{\alpha\mu}} = 0 \quad (2.107)$$

$$A_{\mu\nu} k^\mu \sin(k_\beta x^\beta) = 0. \quad (2.108)$$

Note that this must be true for all x^β , and hence, we have an additional constraint on the four wave-co-vector k_α and the constant tensor $A_{\mu\nu}$ given by

$$A_{\mu\nu} k^\mu = 0. \quad (2.109)$$

(3) Finally, recall that $H_{\mu\nu}$ also needs to satisfy the Einstein's field equations

$$\square^2 H_{\mu\nu} = 0. \quad (2.110)$$

Thus, we have

$$\square^2 H_{\mu\nu} = \partial^\beta \partial_\beta [A_{\mu\nu} \cos(k_\alpha x^\alpha)] = 0. \quad (2.111)$$

$$-A_{\mu\nu} \partial^\beta \sin(k_\sigma x^\sigma) \underbrace{k_\alpha \delta^\alpha_\beta}_{=k_\beta} = 0 \quad (2.112)$$

$$A_{\mu\nu} \cos(k_\rho x^\rho) k_\sigma \delta^{\beta\sigma} k_\beta = 0 \quad (2.113)$$

$$[A_{\mu\nu} \cos(k_\rho x^\rho)] k^\beta k_\beta = 0. \quad (2.114)$$

Again, note that this must be true for all x^ρ , and we do not want $A_{\mu\nu} = 0$, corresponding to the trivial, "no-wave" solution. This leaves us with the constraint

$$k^\beta k_\beta = 0. \quad (2.115)$$

Note that

$$k^\beta k_\beta = -\omega^2 + k^2 = 0 \quad (2.116)$$

$$\omega = k. \quad (2.117)$$

Since wave speed v is related to the wave angular frequency ω and wave-vector k as

$$vk = \omega. \quad (2.118)$$

Hence, the speed of the gravitational waves is given by

$$v = \frac{\omega}{k} = 1 \quad (2.119)$$

which means that gravitational waves travel at the speed of light (since we have set the speed of light $c = 1$).

To further simplify the solutions, we consider the gauge

$$\xi^\alpha = B^\alpha \sin(k_\rho x^\rho), \quad (2.120)$$

where B^α is just another constant four-vector. We have

$$\square^2 \xi_\nu = \partial^\mu \partial_\mu B_\nu \sin(k_\rho x^\rho) \quad (2.121)$$

$$= B_\nu \partial^\mu [\cos(k_\beta x^\beta) k_\rho \delta^\rho_\mu] \quad (2.122)$$

$$= -B_\nu \sin(k_\sigma x^\sigma) k_\beta \delta^{\beta\mu} k_\mu \quad (2.123)$$

$$= -B_\nu \sin(k_\sigma x^\sigma) \underbrace{k^\mu k_\mu}_{=0} \quad (2.124)$$

$$= 0 \quad (2.125)$$

where the last line is obtained based on the constraint placed on k_α from the Einstein's field equations. Recall that under a gauge transformation, the tensor $H_{\mu\nu}$ transforms as

$$\tilde{H}_{\mu\nu} = H_{\mu\nu} - \xi_{\mu,\nu} - \xi_{\nu,\mu} + \eta_{\mu\nu} \square \cdot \vec{\xi}. \quad (2.126)$$

Note that

$$\xi_{\mu,\nu} = \partial_\nu [B_\mu \sin(k_\alpha x^\alpha)] = B_\mu \cos(k_\beta x^\beta) k_\nu, \quad (2.127)$$

and

$$\eta_{\mu\nu} \square \cdot \vec{\xi} = \eta_{\mu\nu} \partial^\alpha \xi_\alpha = \eta_{\mu\nu} B_\alpha \cos(k_\beta x^\beta) k^\alpha. \quad (2.128)$$

Hence we have

$$\tilde{A}_{\mu\nu} \cos(k_\beta x^\beta) = A_{\mu\nu} \cos(k_\beta x^\beta) - B_\mu \cos(k_\beta x^\beta) k_\nu \quad (2.129)$$

$$- B_\nu \cos(k_\beta x^\beta) k_\mu + \eta_{\mu\nu} B_\alpha \cos(k_\beta x^\beta) k^\alpha \quad (2.130)$$

$$\tilde{A}_{\mu\nu} = A_{\mu\nu} - B_\mu k_\nu - B_\nu k_\mu + \eta_{\mu\nu} B_\alpha k^\alpha. \quad (2.131)$$

If we want to eliminate the time-components in the transformed $A_{\mu\nu}$, we require

$$A_{tv} - B_t k_v - B_v k_t + \eta_{tv} B_\alpha k^\alpha = 0. \quad (2.132)$$

We further require that $A_{\mu\nu}$ to be “traceless”:

$$g^{\mu\nu} A_{\mu\nu} \approx \eta^{\mu\nu} A_{\mu\nu} = A^\mu{}_\mu = 0, \quad (2.133)$$

by imposing the constraint

$$g^{\mu\nu} \tilde{A}_{\mu\nu} = g^{\mu\nu} [A_{\mu\nu} - B_\mu k_\nu - B_\nu k_\mu + \eta_{\mu\nu} B_\alpha k^\alpha] = 0 \quad (2.134)$$

$$A^\mu{}_\mu - B_\mu k^\mu - B_\nu k^\nu + 4B_\alpha k^\alpha = 0 \quad (2.135)$$

$$B_\alpha k^\alpha = -\frac{1}{2} A^\mu{}_\mu. \quad (2.136)$$

This simplifies the previous constraint as

$$A_{tv} - B_t k_v - B_v k_t - \frac{1}{2} \eta_{tv} A^\mu{}_\mu = 0. \quad (2.137)$$

We have shown that we can always find a gauge such that (1) the wave solution is traceless, i.e.

$$A^\mu{}_\mu = 0, \quad (2.138)$$

and (2) transverse, i.e.

$$A_{\mu t} = 0 = A_{t\mu}. \quad (2.139)$$

This gauge transformation is known as the traceless-transverse gauge (TT gauge), and we denote the corresponding tensor as $H_{\mu\nu}^{\text{TT}}$.

Recall that the newly defined tensor $H_{\mu\nu}$ is related to the original metric perturbation $h_{\mu\nu}$ as

$$h_{\mu\nu} = H_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} H. \quad (2.140)$$

In the TT gauge, $A_{\mu\nu}$ (and hence $H_{\mu\nu}$) is traceless, and so $H = 0$ which gives

$$h_{\mu\nu}^{\text{TT}} = H_{\mu\nu}^{\text{TT}}. \quad (2.141)$$

Suppose the gravitational wave is travelling along the z direction. This means only the z component of the wave-vector \vec{k} is non-zero. Then, the four wave-co-vector becomes

$$k_\alpha = [-\omega, 0, 0, \omega]. \quad (2.142)$$

(Note that $k_\alpha k^\alpha = 0$.)

In general, we have $H_{\mu\nu} = A_{\mu\nu} \cos(k_\alpha x^\alpha)$. Under the TT gauge with the assumption that the wave is travelling along the z direction, we have

$$\cos(k_\alpha x^\alpha) = \cos(-k^t t + k^x x + k^y y + k^z z) = \cos(kz - \omega t). \quad (2.143)$$

Now, the tensor $A_{\mu\nu}$ in general consists of 16 elements, which can be written explicitly as

$$A_{\mu\nu} = \begin{bmatrix} A_{tt} & A_{tx} & A_{ty} & A_{tz} \\ A_{xt} & A_{xx} & A_{xy} & A_{xz} \\ A_{yt} & A_{yx} & A_{yy} & A_{yz} \\ A_{zt} & A_{zx} & A_{zy} & A_{zz} \end{bmatrix}. \quad (2.144)$$

Note that even without working in the TT-gauge, not all 16 elements in $A_{\mu\nu}$ are independent of each other because $A_{\mu\nu}$ is symmetric, i.e. $A_{\mu\nu} = A_{\nu\mu}$. There are only 10 independent elements without the TT gauge transformation. Working in the TT gauge, we have $A_{t\nu} = 0$. This means that

$$A_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & A_{xz} \\ 0 & A_{xy} & A_{yy} & A_{yz} \\ 0 & A_{xz} & A_{yz} & A_{zz} \end{bmatrix}. \quad (2.145)$$

This leaves us with 6 independent elements. We also have $A^\mu{}_\mu = 0$, which means

$$A_{xx} + A_{yy} + A_{zz} = 0. \quad (2.146)$$

This further reduces our degree of freedom by 1, leaving us with 5 independent elements. Recall that the Lorenz Gauge requires that

$$A_{\mu\nu} k^\mu = 0. \quad (2.147)$$

Imposing this constraint gives

$$A_{t\nu} \omega + A_{z\nu} \omega = 0, \quad (2.148)$$

since only the t and z components are non-zero. Note that we require $A_{t\nu} = 0$. This leaves us with

$$A_{z\nu} \omega = 0. \quad (2.149)$$

This means $A_{zx} = A_{zy} = A_{zz} = 0$. Hence we have

$$A_{\mu\nu} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & A_{yy} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad (2.150)$$

i.e. we are left just 3 independent elements. However, note that we have an additional constraint that $A^\mu{}_\mu = 0$, which leads to

$$A_{xx} + A_{yy} = 0 \quad (2.151)$$

$$A_{xx} = -A_{yy}. \quad (2.152)$$

This leaves us with just 2 independent elements. Hence, we can write the metric perturbation matrix as

$$\begin{aligned} H_{\mu\nu} = A_{\mu\nu} \cos(kz - \omega t) &= \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \cos(kz - \omega t) \\ &= \left(A_{xx} \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} + A_{xy} \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) \cos(kz - \omega t). \end{aligned} \quad (2.153)$$

$$(2.154)$$

We will rename the remaining independent parameters A_{xx} and A_{xy} as A_+ and A_\times respectively, such that

$$H_{\mu\nu} = \left(A_+ \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} + A_\times \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \right) \cos(kz - \omega t). \quad (2.155)$$

The A_+ and A_\times are known as the plus polarization and cross polarization respectively.

Let us first consider we have a particle that is initially at rest. To the leading order, the four-velocity of the particle is given by

$$u^\mu = [1, 0, 0, 0]. \quad (2.156)$$

Now, the geodesic equation

$$\frac{d^2 x^\mu}{d\tau^2} = -\Gamma^\mu_{\alpha\beta} u^\alpha u^\beta, \quad (2.157)$$

similar to the Lagrange equation in classical mechanics, gives us the particle's equation of motion when the gravitational wave passes through. Recall that in terms of the metric, the Christoffel symbol can be written as

$$\Gamma^\mu_{\alpha\beta} = \frac{1}{2} g^{\mu\sigma} \left[g_{\alpha\sigma,\beta} + g_{\sigma\beta,\alpha} - g_{\alpha\beta,\sigma} \right]. \quad (2.158)$$

Since only the u^t component is non-zero, we have

$$\frac{d^2 x^\mu}{d\tau^2} = -\Gamma^\mu_{tt} \underbrace{u^t u^t}_{=1} \quad (2.159)$$

$$= -\frac{1}{2} g^{\mu\sigma} \left[g_{t\sigma,t} + g_{\sigma t,t} - g_{tt,\sigma} \right]. \quad (2.160)$$

Note that $g_{t\sigma} = 0$ except for $g_{tt} = -1$, but differentiating it also gives zero. Hence, the three terms in the bracket vanish, leaving us with

$$\frac{d^2 x^\mu}{d\tau^2} = 0. \quad (2.161)$$

Nevertheless, this does not mean the particle is not moving as a gravitational wave passes through. It simply means that in the TT-gauge, the coordinates of the particle are always the same (i.e. the coordinate system we adopted is a comoving coordinate frame). To visualize the motion of the particle with a gravitational wave passing through, we consider a particle lying on the $z = 0$ plane at some distance R away from the origin with polar angle θ . In particular, we assume

$$\begin{cases} \Delta x = R \cos \theta \\ \Delta y = R \sin \theta \end{cases}. \quad (2.162)$$

With the particle lying on the plane $z = 0$, the metric tensor can be written as

$$g_{\alpha\beta} = \eta_{\alpha\beta} + H_{\alpha\beta} \quad (2.163)$$

$$= \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & A_{xx} & A_{xy} & 0 \\ 0 & A_{xy} & -A_{xx} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \cos[k(0) - \omega t] \quad (2.164)$$

$$= \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 + A_{xx} \cos(\omega t) & A_{xy} \cos(\omega t) & 0 \\ 0 & A_{xy} \cos(\omega t) & 1 - A_{xx} \cos(\omega t) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (2.165)$$

Now, when a gravitational wave passes through, the physical distance between the particle and the origin is given by

$$\Delta s^2 = g_{\alpha\beta} \Delta x^\alpha \Delta x^\beta. \quad (2.166)$$

In the TT gauge, this becomes

$$\Delta s^2 = g_{\alpha\beta} \Delta x^\alpha \Delta x^\beta \quad (2.167)$$

$$= (\Delta x^\alpha)^T g_{\alpha\beta} (\Delta x^\beta) \quad (2.168)$$

$$= \begin{bmatrix} 0 & R \cos \theta & R \sin \theta & 0 \end{bmatrix} \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 + A_{xx} \cos(\omega t) & A_{xy} \cos(\omega t) & 0 \\ 0 & A_{xy} \cos(\omega t) & 1 - A_{xx} \cos(\omega t) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 \\ R \cos \theta \\ R \sin \theta \\ 0 \end{bmatrix} \quad (2.169)$$

$$= R^2 \begin{bmatrix} 0 \\ (1 + A_{xx} \cos(\omega t)) \cos \theta + A_{xy} \sin \theta \cos(\omega t) \\ A_{xy} \cos \theta \cos(\omega t) + (1 - A_{xx} \cos(\omega t)) \sin \theta \\ 0 \end{bmatrix}^T \begin{bmatrix} 0 \\ \cos \theta \\ \sin \theta \\ 0 \end{bmatrix} \quad (2.170)$$

$$= R^2 [\cos^2 \theta + \sin^2 \theta + A_{xx} \cos(\omega t)(\cos^2 \theta - \sin^2 \theta) + 2A_{xy} \cos(\omega t) \sin \theta \cos \theta] \quad (2.171)$$

$$= R^2 [1 + A_{xx} \cos(2\theta) \cos(\omega t) + A_{xy} \sin(2\theta) \cos(\omega t)]. \quad (2.172)$$

Suppose the “cross-polarization” is absent, i.e. $A_{xy} = A_{\times} = 0$, then

$$\Delta s = \sqrt{R^2 [1 + A_+ \cos(2\theta) \cos(\omega t)]} \quad (2.173)$$

$$= R [1 + A_+ \cos(2\theta) \cos(\omega t)]^{1/2} \quad (2.174)$$

$$\approx R \left[1 + \frac{A_+}{2} \cos(2\theta) \cos(\omega t) \right]. \quad (2.175)$$

If we plot the position of the particle at different θ at different times, we will obtain Figure 2.4.1. On the other hand, assuming that the “plus-polarization” is absent, i.e. $A_{xx} = A_+ = 0$, then

$$\Delta s = \sqrt{R^2 [1 + A_+ \sin(2\theta) \cos(\omega t)]} \quad (2.176)$$

$$= R [1 + A_+ \sin(2\theta) \cos(\omega t)]^{1/2} \quad (2.177)$$

$$\approx R \left[1 + \frac{A_+}{2} \sin(2\theta) \cos(\omega t) \right]. \quad (2.178)$$

If we plot the position of the particle at different θ at different times, we will obtain Figure 2.4.2.




Figure 2.4.1: How a ring of test masses on the x - y plane deforms as a gravitational wave with plus-polarization travelling in the z -direction passes by.




Figure 2.4.2: How a ring of test masses on the x - y plane deforms as a gravitational wave with cross-polarization traveling in the z -direction passes by.

2.5 Energy carried by gravitational waves

To understand the amount of energy carried by gravitational waves, we return to the Einstein's field equation:

$$G_{\mu\nu} = 8\pi G T_{\mu\nu}. \quad (2.179)$$

In earlier sections we assumed that $T_{\mu\nu} = 0$, i.e. the wave is travelling in free space. Now, we put the stress-energy term back into the equation. However, this term tells us nothing about the gravitational wave itself because it only encodes information about the system that generates the gravitational waves, e.g. the energy distribution in a binary black hole system. We must consider higher-order terms to extract information about the energy carried by gravitational waves. As a first approximation, we start with second-order terms. To first order, we have

$$G_{\mu\nu}^{(1)} = 8\pi G T_{\mu\nu}, \quad (2.180)$$

with $G_{\mu\nu}^{(1)} = -\frac{1}{2}\Box^2 H_{\mu\nu}$. Promoting this to the second order gives

$$G_{\mu\nu}^{(1)} + G_{\mu\nu}^{(2)} = 8\pi G T_{\mu\nu}. \quad (2.181)$$

In terms of $H_{\mu\nu}$, we can write this as

$$G_{\mu\nu}^{(1)} + G_{\mu\nu}^{(2)} = 8\pi G T_{\mu\nu} \quad (2.182)$$

$$-\frac{1}{2}\square^2 H_{\mu\nu} = 8\pi G T_{\mu\nu} - G_{\mu\nu}^{(2)} \quad (2.183)$$

$$\square^2 H_{\mu\nu} = -16\pi G \left(T_{\mu\nu} - \frac{G_{\mu\nu}^{(2)}}{8\pi G} \right). \quad (2.184)$$

The first term on the right-hand side is just $T_{\mu\nu}$, i.e. the stress-energy tensor for the source of gravitational waves. The second term comes from the geometric side, i.e. it comes directly from the gravitational waves generated. This term encodes information about the energy the gravitational waves carry. By moving this term to the right-hand side of the equation, we are treating it as if it is yet another stress-energy tensor term. This additional term will tell us about the energy of gravitational waves.

It should be noted that the second term itself evaluated at a single point in spacetime is meaningless. Instead, we need to average this term over several gravitational-wave wavelengths. That is, the gravitational-wave energy tensor is defined as

$$T_{\mu\nu}^{\text{GW}} = -\frac{\langle G_{\mu\nu}^{(2)} \rangle}{8\pi G}. \quad (2.185)$$

For a gravitational wave with only plus polarization content, assuming that the gravitational-wave frequency is ω , and that its direction of propagation is $+z$, the metric tensor can be written as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 + A_+ \cos(\omega t - \omega z) & 0 & 0 \\ 0 & 0 & 1 - A_+ \cos(\omega t - \omega z) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (2.186)$$

For simplicity, we write $A_+ \cos(\omega t - \omega z) = h_+(t, z)$ for the time being. To the second order, evaluating the Ricci tensor gives

$$R_{tt} = R_{zz} \approx h_+ \ddot{h}_+ + \frac{1}{2} \dot{h}_+^2, \quad (2.187)$$

and

$$R_{tz} = R_{zt} \approx \dot{h}_{+,z} h_+ + \frac{1}{2} \dot{h}_+ h_{+,z}. \quad (2.188)$$

All other elements in the Ricci tensor are identically 0. Then, the Ricci scalar is given by

$$R \approx \eta^{\mu\nu} R_{\mu\nu} = -R_{tt} + R_{zz} = 0. \quad (2.189)$$

Recall that the Einstein tensor is defined as

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}R. \quad (2.190)$$

The second-order correction term for the Einstein tensor is hence given by

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}R = R_{\mu\nu} = \begin{bmatrix} h_+\ddot{h}_+ + \frac{1}{2}\dot{h}_+^2 & 0 & 0 & \dot{h}_{+,z}h_+ + \frac{1}{2}\dot{h}_+h_{+,z} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \dot{h}_{+,z}h_+ + \frac{1}{2}\dot{h}_+h_{+,z} & 0 & 0 & h_+\ddot{h}_+ + \frac{1}{2}\dot{h}_+^2 \end{bmatrix} \quad (2.191)$$

To find the stress-energy tensor for gravitational waves, we need the wavelength-averaged expression for the Einstein tensor, i.e. $\langle G_{\mu\nu} \rangle$. With

$$h_+ = A_+ \cos(\omega t - \omega z), \quad (2.192)$$

we obtain

$$\ddot{h}_+ = -A_+\omega^2 \cos(\omega t - \omega z), \quad (2.193)$$

and

$$\dot{h}_+ = -A_+\omega \sin(\omega t - \omega z). \quad (2.194)$$

This leads to

$$\langle h_+\ddot{h}_+ + \frac{1}{2}\dot{h}_+^2 \rangle = \langle -A_+^2\omega^2 \cos^2(\omega t - \omega z) + \frac{1}{2}A_+^2\omega^2 \sin^2(\omega t - \omega z) \rangle \quad (2.195)$$

$$= -A_+^2\omega^2 \langle \cos^2(\omega t - \omega z) \rangle + \frac{1}{2}A_+^2\omega^2 \langle \sin^2(\omega t - \omega z) \rangle \quad (2.196)$$

$$= \left(-\frac{1}{2} + \frac{1}{4} \right) A_+^2\omega^2 = -\frac{1}{4}A_+^2\omega^2. \quad (2.197)$$

Note that

$$-\frac{1}{2}\langle \dot{h}_+^2 \rangle = -\frac{1}{2}A_+^2\omega^2 \langle \sin^2(\omega t - \omega z) \rangle = -\frac{1}{4}A_+^2\omega^2. \quad (2.198)$$

This means that for the non-zero diagonal elements in the Einstein tensor, the averaged value will be equivalent to $-\frac{1}{2}\langle\dot{h}_+^2\rangle$, and hence, the tt component of the stress-energy tensor for gravitational waves with only plus-polarization content is

$$T_{tt}^{\text{GW}} = -\frac{G_{tt}^{(2)}}{8\pi G} = \frac{\langle\dot{h}_+^2\rangle}{16\pi G}. \quad (2.199)$$

Similarly, the stress-energy tensor for gravitational waves with only cross-polarization content is given by

$$T_{tt}^{\text{GW}} = \frac{\langle\dot{h}_\times^2\rangle}{16\pi G}. \quad (2.200)$$

For general gravitational waves that have both plus and cross-polarization content, the stress-energy tensor is then given by

$$T_{tt}^{\text{GW}} = \frac{\langle\dot{h}_+^2 + \dot{h}_\times^2\rangle}{16\pi G}. \quad (2.201)$$

Now, the energy density of a gravitational wave is simply the tt component of the stress-energy tensor. Note that

$$h_{\mu\nu}^{\text{TT}} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}. \quad (2.202)$$

Since

$$\dot{h}_{\mu\nu}^{\text{TT}}\dot{h}^{\mu\nu}_{\text{TT}} = (\dot{h}_+)^2 + (\dot{h}_\times)^2 + (\dot{h}_\times)^2 + (-\dot{h}_+)^2 = 2(\dot{h}_+^2 + \dot{h}_\times^2), \quad (2.203)$$

we have

$$T_{tt}^{\text{GW}} = \frac{1}{2} \times \frac{\langle\dot{h}_{\mu\nu}^{\text{TT}}\dot{h}^{\mu\nu}_{\text{TT}}\rangle}{16\pi G} = \frac{\langle\dot{h}_{\mu\nu}^{\text{TT}}\dot{h}^{\mu\nu}_{\text{TT}}\rangle}{32\pi G} = \frac{\langle\dot{h}_+^2 + \dot{h}_\times^2\rangle}{16\pi G}. \quad (2.204)$$

As we have assumed that the gravitational wave is propagating along the z direction, The tz component of the stress-energy tensor gives the flux of gravitational wave energy. The tz component of the Einstein tensor is given by

$$G_{tz} = \dot{h}_{+,z}h_+ + \frac{1}{2}\dot{h}_+h_{+,z}. \quad (2.205)$$

The averaged expression is given by

$$\langle G_{tz} \rangle = \frac{1}{2}\langle\dot{h}_+^2\rangle. \quad (2.206)$$

Generalizing to gravitational waves with mixed polarizations, the expression becomes

$$\langle G_{tz} \rangle = \frac{1}{4} \langle h_{\mu\nu}^{\text{TT}} h_{\text{TT}}^{\mu\nu} \rangle, \quad (2.207)$$

and hence the energy flux of gravitational wave is simply the same as the energy density of the wave itself, apart from an overall sign difference (since it is transporting energy).

2.6 The LIGO detectors: An overview

Brief introduction to the LIGO detectors




Figure 2.6.1: A simplified schematic diagram of the Advanced LIGO gravitational-wave detector. Figure from [341].

A schematic overview of the gravitational wave detector used by LIGO is shown in Figure 2.6.1. Advanced LIGO gravitational-wave detectors are essentially modified Michelson interferometers [341]. A pair of test masses with highly reflective mirror surfaces forms a Fabry-Perot arm cavity along each of the two detector arms. A power recycling mirror is placed between the input laser and the beamsplitter in the lower left part of the figure to enhance the laser power stored in the detector arms to 100 kW. The laser beams travelling along the two detector arms are joined at the beamsplitter and directed to the gravitational-wave readout photodetector, as shown in 2.6.1. We place a signal-recycling mirror between the beam splitter and the gravitational-wave readout photodetector to alter the frequency response of the

interferometer to fluctuations in differential arm length. When gravitational waves pass through the gravitational-wave detectors, it will cause apparent changes in the differential arm lengths that induce power fluctuations in the detectors' gravitational-wave readout port. With appropriate calibration, the signal from the gravitational-wave readout port can be converted into a change in differential arm lengths of the detectors, formally known as strain $h(t)$, in the time domain. Search pipelines can then ingest the calibrated strain data to search for possible gravitational-wave signals. In Chapter 3, we will discuss how data from LIGO detectors is calibrated.

Antenna pattern associated with LIGO's L-shaped detectors

We will now derive the antenna pattern associated with the geometry of the gravitational-wave detectors. For LIGO detectors, it corresponds to an L-shaped ground-based detector. Consider a gravitational wave travelling along the direction \hat{n} , the corresponding strain tensor is then given by

$$h_{jk}(t, \vec{r}) = \sum_{A=+, \times} e_{jk}^A(\hat{n}) \int h_A(f) \exp \left[2\pi i f \left(t - \frac{\hat{n} \cdot \vec{r}}{c} \right) \right], \quad (2.208)$$

where A represents the polarization ("plus" and "cross" polarization). The Latin alphabet indices j and k represent the three spatial dimensions (i.e. excluding the time dimension)². The integral is an inverse Fourier transform of the strain amplitude of gravitational wave evaluated at retarded time $t - \frac{\hat{n} \cdot \vec{r}}{c}$. $e_{jk}^A(\hat{n})$ are the polarization tensors for gravitational waves in the TT gauge, specifying the test mass ring motion:

$$\begin{cases} e_{jk}^+ = \hat{u}_j \hat{u}_k - \hat{v}_j \hat{v}_k \\ e_{jk}^\times = \hat{u}_j \hat{v}_k + \hat{v}_j \hat{u}_k \end{cases}. \quad (2.209)$$

where \hat{u} and \hat{v} are unit vectors orthogonal to \hat{n} and to each other. Without loss of generality, we set $\vec{r} = 0$ as the position of the LIGO detector.

When the gravitational wavelength $\lambda_{GW} = \frac{1}{f}$ is much longer than the detector's typical size, we can neglect the second term in the retarded time, i.e.

$$e^{-i2\pi f(t - \vec{n} \cdot \vec{r})} \approx e^{-i2\pi f t}. \quad (2.210)$$

² In the rest of this thesis, Greek alphabet indices represent time and spatial dimensions (i.e. 4-dimension) where Latin alphabet indices represent spatial dimensions only.

With this, the tensor strain can be simplified as

$$h_{jk}(t) = \sum_{A=+, \times} e_{jk}^A(\hat{n}) \int_{-\infty}^{\infty} df \tilde{h}_A(f) e^{-i2\pi f(t-\hat{n}\cdot\vec{r})} df \quad (2.211)$$

$$\approx \sum_{A=+, \times} e_{jk}^A(\hat{n}) \int_{-\infty}^{\infty} df \tilde{h}_A(f) e^{-i2\pi f t} df \quad (2.212)$$

$$= \sum_{A=+, \times} e_{jk}^A(\hat{n}) h_A(t) \quad (2.213)$$

where we make use of the fact that the integral in the second line is simply the “inverse Fourier transform” of the strain amplitude in frequency space $\tilde{h}_A(f)$.

Recall that the input of the detector is given by

$$h(t) = D^{jk} h_{ij}(t), \quad (2.214)$$

where D^{jk} is a tensor describing the detector’s geometry. We define the detector pattern functions as

$$F_A(\hat{n}) = D^{jk} e_{jk}^A(\hat{n}). \quad (2.215)$$

In terms of spherical coordinates (r, θ, ϕ) , the general expression for the directional unit vector \hat{n} can be expressed as

$$\hat{n} = \sin \theta \cos \phi \hat{x} + \sin \theta \sin \phi \hat{y} + \cos \theta \hat{z}. \quad (2.216)$$

In terms of the antenna pattern functions and the two polarizations of the gravitational wave h_+ and h_\times , we can write

$$h(t) = D^{jk} h_{jk}(t) = D^{jk} h_{jk}(t) \sum_{A=+, \times} e_{ij}^A(\hat{n}) h_A(t) \quad (2.217)$$

$$= \sum_{A=+, \times} [D^{jk} h_{jk}(t) e_{jk}^A] h_A(t) \quad (2.218)$$

$$= \sum_{A=+, \times} F_A h_A(t) \quad (2.219)$$

$$= F_+ h_+ + F_\times h_\times. \quad (2.220)$$

We have assumed above that we have chosen a system of axes (\hat{u}, \hat{v}) in the plane orthogonal to the propagation direction \hat{n} of the wave, with respect to which the polarizations h_+ and h_\times are defined. In the case where the system is rotated along the \hat{n} axis counter-clockwise by an angle ψ . Take u and v as the x and y axes

respectively. The rotational matrix \mathbf{R} that represents the rotational transformation from the $u - v$ plane to the $u' - v'$ plane is given by

$$\mathbf{R} = \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix}. \quad (2.221)$$

From this, we have

$$\begin{bmatrix} u' \\ v' \end{bmatrix} = \mathbf{R} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} u \cos \psi - v \sin \psi \\ u \sin \psi + v \cos \psi \end{bmatrix}. \quad (2.222)$$

The original TT gauge expression for h_{jk}^{TT} is given by

$$h_{jk}^{TT} = \begin{bmatrix} h_+ & h_\times \\ h_\times & -h_+ \end{bmatrix} \quad (2.223)$$

where I have neglected the t and z components. By transforming the strain tensor h_{jk}^{TT} with the rotational matrix, we have

$$h_{j'k'}^{TT} = \begin{bmatrix} h'_+ & h'_\times \\ h'_\times & -h'_+ \end{bmatrix} \quad (2.224)$$

$$= \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} \begin{bmatrix} h_+ & h_\times \\ h_\times & -h_+ \end{bmatrix} \begin{bmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{bmatrix} \quad (2.225)$$

$$= \begin{bmatrix} h_+ \cos \psi - h_\times \sin \psi & h_\times \cos \psi + h_+ \sin \psi \\ h_+ \sin \psi + h_\times \cos \psi & h_\times \sin \psi - h_+ \cos \psi \end{bmatrix} \begin{bmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{bmatrix} \quad (2.226)$$

$$= \begin{bmatrix} h_+(\cos^2 \psi - \sin^2 \psi) - h_\times(2 \sin \psi \cos \psi) & h_+(2 \sin \psi \cos \psi) + h_\times(\cos^2 \psi - \sin^2 \psi) \\ h_+(2 \sin \psi \cos \psi) + h_\times(\cos^2 \psi - \sin^2 \psi) & -[h_+(\cos^2 \psi - \sin^2 \psi) - h_\times(2 \sin \psi \cos \psi)] \end{bmatrix} \quad (2.227)$$

$$= \begin{bmatrix} h_+ \cos(2\psi) - h_\times \sin(2\psi) & h_+ \sin(2\psi) + h_\times \cos(2\psi) \\ h_+ \sin(2\psi) + h_\times \cos(2\psi) & -[h_+ \cos(2\psi) - h_\times \sin(2\psi)] \end{bmatrix} \quad (2.228)$$

and hence

$$\begin{cases} h'_+ = h_+ \cos(2\psi) - h_\times \sin(2\psi) \\ h'_\times = h_+ \sin(2\psi) + h_\times \cos(2\psi) \end{cases}. \quad (2.229)$$

Then, in the rotated frame, the plus polarization tensor can be written as

$$e_{jk}^{+'} = \hat{u}'_j \hat{u}'_k - \hat{v}'_j \hat{v}'_k \quad (2.230)$$

$$= (\hat{u}_j \cos \psi - \hat{v}_j \sin \psi)(\hat{u}_k \cos \psi - \hat{v}_k \sin \psi) - (\hat{u}_j \sin \psi + \hat{v}_j \cos \psi)(\hat{u}_k \sin \psi + \hat{v}_k \cos \psi) \quad (2.231)$$

$$= [\hat{u}_j \hat{u}_k - \hat{v}_j \hat{v}_k](\cos^2 \psi - \sin^2 \psi) - [\hat{u}_j \hat{v}_k + \hat{v}_j \hat{u}_k](2 \sin \psi \cos \psi) \quad (2.232)$$

$$= \cos(2\psi) e_{jk}^+ - \sin(2\psi) e_{jk}^\times, \quad (2.233)$$

and similarly, the cross-polarization tensor can be written as

$$e_{jk}^{\times'} = \hat{u}'_j \hat{v}'_k + \hat{v}'_j \hat{u}'_k \quad (2.234)$$

$$= (\hat{u}_j \cos \psi - \hat{v}_j \sin \psi)(\hat{u}_k \sin \psi + \hat{v}_k \cos \psi) + (\hat{u}_j \cos \psi - \hat{v}_j \sin \psi)(\hat{u}_k \cos \psi - \hat{v}_k \sin \psi) \quad (2.235)$$

$$= [\hat{u}_j \hat{u}_k - \hat{v}_j \hat{v}_k](2 \sin \psi \cos \psi) + [\hat{u}_j \hat{v}_k + \hat{v}_j \hat{u}_k](\cos^2 \psi - \sin^2 \psi) \quad (2.236)$$

$$= \sin(2\psi)e_{jk}^+ + \cos(2\psi)e_{jk}^\times. \quad (2.237)$$

With the detector antenna pattern functions defined as

$$F_A(\hat{n}) = D^{jk} e_{jk}^A(\hat{n}), \quad (2.238)$$

since D^{ij} depends only on the geometry of the detector, it should remain unchanged under the rotation. In the new rotated system, we have

$$\begin{cases} F'_+(\hat{n}, \psi) = F_+(\hat{n}, 0) \cos(2\psi) - F_\times(\hat{n}, 0) \sin(2\psi) \\ F'_\times(\hat{n}, \psi) = F_+(\hat{n}, 0) \sin(2\psi) + F_\times(\hat{n}, 0) \cos(2\psi) \end{cases}. \quad (2.239)$$

Now consider Figure 2.6.2, which shows the gravitational-wave source frame (left) and the LIGO detector frame (right). The polarizations h_+ and h_\times are defined with respect to the $x' - y'$ frame. Therefore, in the source frame, the gravitational wave strain tensor has the form (note that I have ignored the time component)

$$h'_{jk} = \begin{bmatrix} h'_{xx} & h'_{xy} & h'_{xz} \\ h'_{yx} & h'_{yy} & h'_{yz} \\ h'_{zx} & h'_{zy} & h'_{zz} \end{bmatrix} = \begin{bmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (2.240)$$

To transform h'_{jk} into the detector frame, we need to do two rotations, respectively along the y axis and z axis. The rotational matrix about the y axis is given by

$$\mathbf{R}_y = \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}. \quad (2.241)$$

and that about the z axis is given by

$$\mathbf{R}_z = \begin{bmatrix} \cos \phi & \sin \phi & 0 \\ -\sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (2.242)$$




Figure 2.6.2: A schematic diagram illustrating gravitational-wave source frame and detector frame.

Therefore, we have

$$\mathbf{R} = \mathbf{R}_z \mathbf{R}_y = \begin{bmatrix} \cos \phi & \sin \phi & 0 \\ -\sin \phi & \cos \phi & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix} \quad (2.243)$$

$$= \begin{bmatrix} \cos \phi \cos \theta & \sin \phi & \cos \phi \sin \theta \\ -\sin \phi \cos \theta & \cos \phi & -\sin \phi \sin \theta \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}. \quad (2.244)$$

In terms of h_+ , h_\times , θ and ϕ , we have

$$h_{jk} = \mathbf{R} h'_{jk} \mathbf{R}^T \quad (2.245)$$

$$= \begin{bmatrix} \cos \phi \cos \theta & \sin \phi & \cos \phi \sin \theta \\ -\sin \phi \cos \theta & \cos \phi & -\sin \phi \sin \theta \\ -\sin \theta & 0 & \cos \theta \end{bmatrix} \begin{bmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \cos \phi \cos \theta & -\sin \phi \cos \theta & -\sin \theta \\ \sin \phi & \cos \phi & 0 \\ \cos \phi \sin \theta & -\sin \phi \sin \theta & \cos \theta \end{bmatrix} \quad (2.246)$$

$$= \begin{bmatrix} h_+ \cos \phi \cos \theta + h_\times \sin \phi & h_\times \cos \phi \cos \theta - h_+ \sin \phi & 0 \\ -h_+ \sin \phi \cos \theta + h_\times \cos \phi & -h_\times \sin \phi \cos \theta - h_+ \cos \phi & 0 \\ -h_+ \sin \theta & -h_\times \sin \theta & 0 \end{bmatrix} \begin{bmatrix} \cos \phi \cos \theta & -\sin \phi \cos \theta & -\sin \theta \\ \sin \phi & \cos \phi & 0 \\ \cos \phi \sin \theta & -\sin \phi \sin \theta & \cos \theta \end{bmatrix} \quad (2.247)$$

$$= \begin{bmatrix} h_+(\cos^2 \phi \cos^2 \theta - \sin^2 \phi) + h_\times \cos \theta \sin(2\phi) & -h_+(\cos^2 \theta + 1) \sin \phi \cos \phi + h_\times \cos \theta \cos(2\phi) & -\frac{1}{2} h_+ \cos \phi \sin(2\theta) - h_\times \sin \phi \sin \theta \\ -h_+(\cos^2 \theta + 1) \sin \phi \cos \phi + h_\times \cos \theta \cos(2\phi) & h_+(\sin^2 \phi \cos^2 \theta - \cos^2 \phi) - h_\times \cos \theta \sin(2\phi) & \frac{1}{2} h_+ \sin \phi \sin(2\theta) - h_\times \cos \phi \sin \theta \\ -\frac{1}{2} h_+ \cos \phi \sin(2\theta) - h_\times \sin \phi \sin \theta & \frac{1}{2} h_+ \sin \phi \sin(2\theta) - h_\times \cos \phi \sin \theta & -h_+ \sin^2 \theta \end{bmatrix}. \quad (2.248)$$

To compute the phase shift in a LIGO detector, defined by $\frac{1}{2}(\partial_{tt} h_{xx} - \partial_{tt} h_{yy})$, we need to write the tensor D_{jk} as

$$D_{jk} = \frac{1}{2}(\hat{x}_j \hat{x}_k - \hat{y}_j \hat{y}_k). \quad (2.249)$$

The strain is given by

$$h(t) = \frac{1}{2}(h_{xx} - h_{yy}) = F_+ h_+ + F_\times h_\times. \quad (2.250)$$

Since

$$\frac{1}{2}(h_{xx} - h_{yy}) = \frac{1}{2}\{h_+[\cos^2 \theta(\cos^2 \phi - \sin^2 \phi) + (\cos^2 \phi - \sin^2 \phi)] + 2h_\times \cos \theta \sin(2\phi)\} \quad (2.251)$$

$$= \frac{1}{2}h_+ \cos(2\phi)(\cos^2 \theta + 1) + h_\times \cos \theta \sin(2\phi), \quad (2.252)$$

we obtain the antenna pattern functions as

$$\begin{cases} F_+ = \frac{1}{2} \cos(2\phi)(\cos^2 \theta + 1) \\ F_\times = \cos \theta \sin(2\phi) \end{cases}. \quad (2.253)$$

With these derived expressions, we can now evaluate the input of the detector given by

$$h(t) = D^{jk} h_{jk}(t), \quad (2.254)$$

with D^{jk} being a tensor describing the geometry of the detector. It is important to note that all the expressions defined here are evaluated in the celestial coordinate frame.

Chapter 3

DATA ANALYSIS OF GRAVITATIONAL WAVES FROM COMPACT BINARY COALESCENCES - AN OVERVIEW

3.1 Introduction

In this chapter, I will give an overview of the LIGO-Virgo-KAGRA (LVK) collaboration's framework to search for gravitational waves. Some of the known sources of gravitational waves include: (1) gravitational waves from compact binary coalescences (CBCs), (2) gravitational-wave bursts that may originate from transients like supernovae, (3) continuous gravitational waves from sources emitting gravitational waves continuously at quasi-constant frequency, for example, spinning asymmetric neutron stars, and (4) stochastic gravitational-wave background that consists of a collection of gravitational waves that are not well-localized. In this thesis, however, we will only be focusing on gravitational waves from compact binary coalescences as LIGO has only confidently detected gravitational waves from CBCs at the time of writing.

3.2 Overview of current LVK's framework to search for gravitational waves from compact binary coalescences

There are several key steps in current LVK's framework to search for and analyze gravitational waves from compact binary coalescences: (1) Calibrating detector data, (2) Detector characterization and noise mitigation, (3) Searching for significant possible gravitational-wave candidates, and (4) Performing parameter estimation for possible gravitational-wave candidates.

Calibrating detector data

Before searching for possible gravitational-wave signals, the detector's data must first be calibrated and have their data quality checked, a process more commonly known as "strain calibration"¹

When a typical gravitational wave passes through a detector, the detector arm lengths will only change by a small amount ΔL on the order of $O(10^{19})\text{m}$. The corresponding gravitational wave strain is $h = \frac{\Delta L}{L}$, where L is the detector arm

¹ I will not pretend to be an expert because I do not actively work on strain calibration. This will be something I am going to explore in the future. Here, I will try my best to give a brief introduction.

length (i.e. 4km) (c.f. Equation 2.250 in Chapter 3). To measure the change in arm lengths to that precision, the calibration of the detector data must also attain at least the same precision. We calibrate the detector data carefully using **feedback control loops** [7]. Feedback control loops (see Figure 3.2.1) are used to maintain

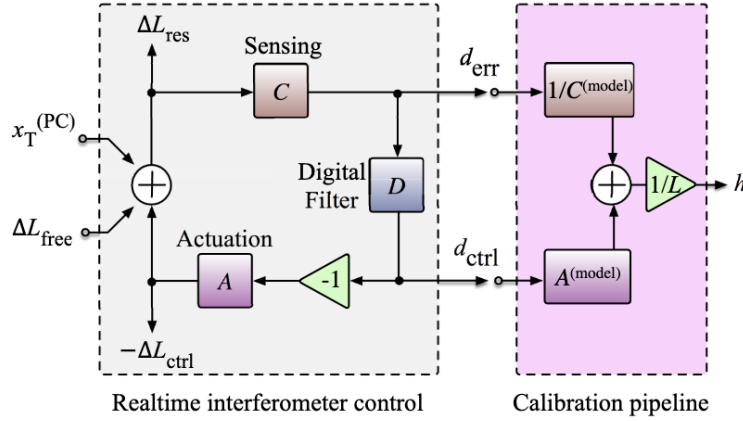


Figure 3.2.1: A block diagram of the feedback control loops used in the LIGO detectors. Figure taken from [341].

the gravitational-wave detectors in a quasi-stationary state so that they can detect gravitational waves. One of these control loops is called the differential arm (DARM) length feedback loop [7]. The DARM length feedback loop, subdivided into sensing and actuation functions, is primarily used to keep the detector arm length difference steady. The sensing function transforms the differential displacement of the test mass to a digitized signal. This digitized signal represents the laser power fluctuation at the gravitational-wave readout port. The actuation function, on the other hand, can control the detector differential arm length by actuating on the quadruple suspension system for any of the arm cavity test masses. Details of the two functions can be found at [7], which will not be repeated here.

For the Advanced LIGO detectors, we use photon calibrators, commonly known as “Pcals”, as the primary calibration tool [207]. The working principle of photon calibrators is as follows: First, auxiliary, power-modulated laser beams are impinged onto a suspended test mass. Because of photon radiation pressure, as the test mass’ mirror surface reflects the photons, the test mass will recoil accordingly. Since the laser beams are modulated, the forces on the test mass are periodic and are directly proportional to the amplitude of the laser power modulation, the test mass’ position will hence undergo modulation together with the length of the detector arm cavity. Figure 3.2.2 illustrates how photon calibrators imping laser beams

on a suspended test mass. We can also use the photon calibrators to produce




Figure 3.2.2: An illustration of how photon calibrators imping laser beams on a suspended test mass. Figure taken from [207].

tracked and known movements to the test masses. This produces “calibration lines” at known frequencies to the detector data. We can use the calibration lines to calibrate the output signal and track slow temporal variation [207]. Calibration lines are excitations induced using the photon calibrators at certain known nominal frequencies with known amplitudes (See Table III of [207]). The use of calibration lines allows one to apply corrections for slow temporal variations and hence improve the accuracy of data calibration. In Figure 4.2.2, two examples of calibration lines at frequencies 36.7 Hz and 331.9 Hz can be seen.

Interested readers are asked to refer to [7, 207] for more information about strain calibration.

Detector characterization and mitigation

While Advanced LIGO detectors have employed various technologies to reduce noise from non-astronomical sources, the complicatedness of the detectors also introduces “instrumental noise” to the data, affecting the data quality and our search sensitivity for gravitational waves. It is, therefore, important to identify, diagnose, and mitigate detector noise that may affect the quality of gravitational-wave data. These procedures are known collectively as **detector characterization** [130].

There are two major instrumental noise types: Transient and persistent noise. Noise transients, more commonly known as “glitches”, which are short-duration bursts of

noise in the data. In contrast, persistent noise is long-term noise residing in the data. The ability to correctly identify and characterize instrumental noise can help us to locate and possibly remove known noise from the data (i.e. data mitigation), providing cleaner data for follow-up analyses to search for and identify gravitational waves from the data. Detailed procedures for identifying and mitigating instrumental noise that is characterize-able in the data can be found in [130, 131].

Because of the Advanced LIGO detectors’ complexity and the wide range of possible noise sources, not all noise in the data can be accurately characterized and mitigated. Data with quality issues that cannot be mitigated will be “vetoed”, and will not be analyzed by analysis pipelines. Identifying and vetoing problematic times of data is particularly important to searches for gravitational waves from compact binary coalescences. This is because search pipelines must be capable of separating bursts of excessive power caused by real gravitational-wave signals, from glitches caused by instrumental noise [130]. LIGO also employs pipelines to help identify times with high glitch rates. For example, Omicron [303] is used to identify and gather all glitch times [130]. A machine-learning-based pipeline, iDQ [149, 176], is also used to find possible correlations with auxiliary witnesses in gravitational-wave data. iDQ aims to automatically identify potential glitches in the gravitational-wave data and provide information about the times in the data up to sub-second intervals likely to contain a glitch in low latency. Search pipelines will incorporate information provided by the aforementioned tools about the data quality and consider them when assigning ranking statistics to triggers found during times with a high probability of corruption by glitches (See also [176]). Later in Chapter 5, we will introduce a proposed method to incorporate data quality information output from the iDQ pipeline into the GstLAL search pipeline.

Since detector characterization is not the main focus of this thesis, readers are asked to refer to [130, 131] for more details.

Searching for possible gravitational-wave signals from the data

Calibrated data are then ingested by “search pipelines” to search for possible candidates of gravitational waves in low-latency. In this thesis, we are mainly focusing on the search for gravitational waves from compact binary coalescences. However, readers are reminded that other sources of gravitational waves are possibly detectable by ground-based gravitational-wave detectors, including continuous gravitational waves, gravitational-wave bursts, and stochastic gravitational-wave background.

Gravitational waves from compact binary coalescences have well-modelled waveforms that typically depend on 15 parameters.




Figure 3.2.3: The 15 parameters that govern the waveforms of gravitational waves from compact binary coalescences.

- m_1, m_2 : Masses of the two components,
- $\vec{\chi}_1, \vec{\chi}_2$: The three-dimensional (dimensionless) spins of the two components,
- ι : Inclination of the orbital plane of the compact binary with respect to the line of sight,
- ψ : The orientation of the polarization vector with respect to the detector,
- D_L : Luminosity distance to the source,
- α, δ : The right ascension and declination of the compact binary,
- t_c, ϕ_c : The coalescence time and phase.

These parameters are also illustrated in Figure 3.2.3. For gravitational waves from compact binaries that involve a neutron star, additional parameters known as the “tidal deformability” Λ_i , corresponding to effects related to the matter and internal structure of neutron stars, are included. In some analyses, the eccentricity of the orbit of the compact binaries is also included as one of the parameters.

There are two major types of search pipelines to search for possible gravitational waves from compact binary coalescences, namely (1) Modelled (matched-filtering-based) searches, and (2) Unmodelled searches.

Modelled (matched-filtering-based) search pipelines, including GstLAL and pyCBC, rely on the fact that gravitational waves from compact binary coalescences have well-modeled waveforms. Given a region in the parameter space (e.g. in the component masses m_1 - m_2 space) where gravitational waves are within the frequency bands that LIGO detectors are sensitive to, we can precompute a set of gravitational waveforms (which we call “template” waveforms) within this region labelled by their respective parameters. This set of precomputed template waveforms is stored collectively in a “template bank”. We then slide each of these templates across the data, and at each point, we compute the correlation between the template and data. The correlation amplitude is known as the signal-to-noise ratio (SNR). This is a process known as “matched-filtering” [43, 189]. Suppose the SNR exceeds a certain pre-determined threshold at a given time instant. In that case, we record the time instant and label it as a possible candidate gravitational-wave signal (also known as a “trigger”). Modelled (matched-filtering-based) search pipelines will then rank these triggers according to some internally assigned ranking statistics from most likely to least likely to be an actual gravitational-wave signal. A ranked list of candidates is then released from the search pipeline for follow-up analysis. Details of matched-filtering-based pipelines will be further discussed in Chapter 4.

In contrast, unmodelled searches like the Coherent WaveBurst (CWB) pipeline [144] aims to identify gravitational waves with minimal assumptions on the signal morphology, i.e., there are no assumed waveform models. For example, the CWB pipeline looks for coherent (i.e., coincident in time) burst signals in the time-frequency domain for data across a network of detectors. Coherent triggers are subsequently assigned a likelihood ratio that compares the probability of the trigger associated with a gravitational-wave signal to that of the trigger originating from noise. If the trigger’s likelihood ratio exceeds a predetermined threshold, it will be flagged by the pipeline as a candidate gravitational wave and passed on for follow-up analyses [215].

While modelled and unmodelled searches have been used to search for gravitational waves, modelled searches, in general, are more effective in uncovering gravitational-wave signals. Some possible reasons are: (1) Modelled (matched-filtering-based) searches can maximize the signal-to-noise ratio (SNR) when the observed data is matched-filtered against a bank of template waveforms by design. If noise in the data is stationary and Gaussian, it has been shown that matched-filtered SNR is the optimal detection statistic to search for possible signals in the data [43].

(2) Unmodelled searches generally require signals to be seen in more than one detector (coherent signals). However, it is not uncommon for detectors in a detector network to be “down” (not in a state where it can collect scientific data) temporarily, leaving only a single detector collecting data. Alternatively, due to the geometry of the detectors, there exists “blind spots” in the sky at any given time from which gravitational waves are not detectable by a detector. In those scenarios, unmodelled searches that rely on the coherent detection of signals across detectors can potentially lose signals. Modelled (matched-filtering-based) searches, on the other hand, can find signals even if they are only detected in one detector ².

In the rest of this thesis, we will focus on modelled (matched-filtering-based) search pipelines.

Bayesian Parameter Estimation for significant events

As a next step, Bayesian parameter estimation is run on significant gravitational-wave candidates output from search pipelines to infer the sources’ astrophysical properties (e.g. the source component masses and spins, c.f. Figure 3.2.3). Given the gravitational-wave data d containing a gravitational-wave signal s , we want to infer the signal’s source parameters $\vec{\theta}$. Formally, we want to determine the posterior $p(\vec{\theta}|d)$ [341], i.e. the probability that the data contains a gravitational-wave signal with source parameters $\vec{\theta}$ given the data d . To evaluate this posterior, we make use of Bayes’ theorem and write the posterior as

$$p(\vec{\theta}|d) = \frac{\mathcal{L}(d|\vec{\theta})\pi(\vec{\theta})}{\mathcal{Z}}, \quad (3.1)$$

where $\mathcal{L}(d|\vec{\theta})$ is the “likelihood” that gives the probability of obtaining the data d if a gravitational-wave signal with source parameters $\vec{\theta}$ is present in the data; $\pi(\vec{\theta})$ is the “prior” distribution for the source parameters $\vec{\theta}$, and

$$\mathcal{Z} = \int \mathcal{L}(d|\vec{\theta})\pi(\vec{\theta})d\vec{\theta} \quad (3.2)$$

is the “evidence”. In this case, the evidence is a normalization constant to the posterior. However, it should be noted that if we are instead performing Bayesian model selection analysis, the “evidence” \mathcal{Z} will become important when evaluating the Bayesian odds [341]. For instance, in Chapter 6, 8, 9 and 10, we perform Bayesian model selection analysis to determine how likely a set of gravitational-wave signals are strongly-lensed counterparts to each other, originating from the

² This requires the signal to have a sufficiently high amplitude in general.

same source, as compared to being independent signals. In such analyses, the “evidence” is no longer a mere normalization constant, but is crucial for evaluating the Bayesian odds ratio (See also [235]).

For most Bayesian parameter estimation analyses like Bilby [49], a Gaussian noise likelihood, defined as

$$\mathcal{L} = \frac{1}{2\pi\sigma^2} \exp \left[-\frac{|d - h(\vec{\theta})|^2}{2\sigma^2} \right], \quad (3.3)$$

where $h(\vec{\theta})$ is a template waveform with parameters $\vec{\theta}$, σ is the detector noise (assuming Gaussian noise), and π is the usual mathematical constant (not the prior), is adopted.

Once the prior and likelihood are set, one can adopt a stochastic sampler, e.g. Markov Chain Monte Carlo (MCMC) based samplers, to explore the parameter space and map out the posterior. Samplers employed by the Bilby parameter estimation pipeline [49] generate a list of posterior samples $\{\vec{\theta}\}$ drawn from the posterior probability distribution such that the number of samples is proportional to the posterior distribution $p(\vec{\theta}|d)$ on the interval $(\vec{\theta}, \vec{\theta} + d\vec{\theta})$ [341].

These posterior samples generated for each candidate gravitational-wave event, while being a by-product of the parameter estimation analysis, are crucial ingredients for other follow-up analyses to be conducted on the gravitational-wave events. For instance, in Chapter 7, posterior samples of confirmed superthreshold gravitational-wave events are used to generate lensed injections for the targeted search for possible sub-threshold lensed gravitational waves.

Details on the exact Bayesian parameter estimation analysis can be found in [49, 341].

Chapter 4

THE GSTLAL SEARCH PIPELINE - AN OVERVIEW

4.1 Top-level introduction to matched-filtering-based search pipelines

Here, we give a general overview of matched-filtering-based search pipelines for gravitational waves. For the rest of this thesis, matched-filtering pipelines are assumed to be targeting only gravitational waves originating from compact binary coalescences, but readers are reminded that these pipelines can potentially be used to search for gravitational waves from other sources (e.g. bursts).

Matched-filtering-based pipelines exploit the fact that gravitational waveforms from compact binary coalescences are well-modelled. Focusing on a particular region in the gravitational-wave parameter space, a collection of waveforms $\{\tilde{h}_i(f)$ (labelled by their intrinsic parameters like component masses $m_{1,2}$ and dimensionless spins $\vec{\chi}_{1,2}$) are prepared and filtered against sets of data $\{d(t)\}$ collected from the participating detectors in an observing run. Assuming aligned spins with the stationary phase approximation, the plus and cross-polarization components of a gravitational waveform can be written in the time domain as

$$\begin{cases} h_+(t) = \frac{1+\cos^2\iota}{2} \left(\frac{2GM}{c^2 D} \right) \left(\frac{2G\mu}{c^2 r(t)} \right) \cos[2(\omega(t) - \phi_c)] \\ h_x(t) = \cos\iota \left(\frac{2GM}{c^2 D} \right) \left(\frac{2G\mu}{c^2 r(t)} \right) \cos[2(\omega(t) - \phi_c)] \end{cases}, \quad (4.1)$$

where ι is the inclination of the system, G is the gravitational constant, c is the speed of light, $M = m_1 + m_2$ is the total mass of the system, μ is the reduced mass of the system, ω is the orbital angular frequency, D is the distance to the source's center, r is the source's separation, and ϕ_c is the coalescence phase. Under this assumption, the two polarizations differ from each other only by a proportionality constant plus a $\frac{\pi}{4}$ phase shift. Note that here we are only considering the fundamental $2-2$ mode, and higher-order modes are neglected. In the future, the inclusion of higher-order modes may be implemented to further improve our search sensitivity [359].

Mathematically, we can write the complex matched-filtering output $z_i(t)$ of the data

$d(t)$ against a given template $\tilde{h}_i(f)$ as

$$z_i(t) \equiv x_i(t) + iy_i(t) = 4 \int_0^\infty \frac{\tilde{d}(f)\tilde{h}_i^*(f)}{S_n(f)} e^{2i\pi ft} df \quad (4.2)$$

$$= 4 \int_0^\infty \frac{\tilde{d}(f)}{\sqrt{S_n(f)}} \cdot \frac{\tilde{h}_i^*(f)}{\sqrt{S_n(f)}} e^{2i\pi ft} df, \quad (4.3)$$

where $S_n(f)$ is the one-sided noise power spectral density (PSD), and the integral in the last line transforms the matched-filtering output from the frequency domain to the time domain via Fourier transform. x_i and y_i correspond to a gravitational wave signal with coalescence phase ϕ_c and $\phi_c + \frac{\pi}{4}$ respectively. The final line of Equation 4.2 explicitly shows that both the data $\tilde{d}(f)$ and the template waveform $\tilde{h}_i(f)$ are whitened by the square root of the noise PSD, i.e. the amplitude spectral density (ASD).

Pipelines will record times when the magnitude of the matched-filtering output. When the output is suitably normalized[43], it is formally known as the signal-to-noise ratio SNR ρ , exceeds a predetermined threshold together with the SNR at those times and the template parameters as a candidate event (also known as “trigger”) for each template waveform. For GstLAL, an additional statistic, the “auto-correlation based consistency test value” ξ^2 , is calculated for each trigger. ξ^2 evaluates the difference between the expected SNR timeseries for a template associated with a trigger (also known as the “auto-correlation”, hence auto-correlation-based) to the actual SNR timeseries. The larger ξ^2 , the more different the data is compared to the template waveform. ξ^2 will be further discussed in Section 4.2. PyCBC has a similar signal consistency test value (see Equation 14 of [351] for example).

Pipelines will then estimate the significance (using a ranking statistic) of the collection of triggers they obtained. PyCBC previously used a ranking statistic called “re-weighted SNR” $\tilde{\rho}$ defined by (See also Equation 5 of [277])

$$\tilde{\rho} = \begin{cases} \rho, & \chi^2 \leq 1 \\ \frac{\rho}{\left[\frac{1}{2} + (\chi^2)^3\right]^{\frac{1}{6}}}, & \chi^2 > 1 \end{cases} \quad (4.4)$$

GstLAL assigns an alternative ranking statistic, the “likelihood ratio”

$$\mathcal{L} = \frac{\mathcal{L}(\text{signal})}{\mathcal{L}(\text{noise})}. \quad (4.5)$$

\mathcal{L} is the ratio of the probabilities of obtaining a trigger (with the recorded SNR ρ , ξ^2 , template parameter $\vec{\theta}$, and other detection statistics) under the signal hypothesis

(i.e. assuming that the trigger is a gravitational-wave signal) and that under the noise hypothesis (i.e. assuming that the trigger is noise). Starting from the third observing run, PyCBC also started to use the “likelihood ratio” as the ranking statistic. This will be further discussed in the next section.

A noise background estimation is required to evaluate the probability under the noise hypothesis. Different search pipelines have different approaches to model the noise background. For example, pyCBC estimates the noise background by performing multiple timeslides of data [277, 279]: In each iteration, pyCBC adds a time offset to the data in different detectors that is larger than the maximum light travel time across the detectors, perform matched-filtering on the time-slided data and collect triggers that are found coincidentally in multiple detectors. These coincident triggers must not correspond to real gravitational waves but noise because of the unphysical time offset added. Hence, they can be used to construct one realization of the noise background. An ensemble of these estimates is then used to generate a representative noise background estimation. On the other hand, GstLAL uses non-coincidental single detector triggers to estimate the noise background (See Section 4.2).

Finally, search pipelines estimate the False-Alarm-Probability (FAP)¹ for each trigger. FAP is the probability that noise can produce a trigger with a ranking statistic larger or equal to that of the trigger under consideration. Intuitively, the lower the FAP, the more likely a candidate is a gravitational wave. A related statistic known as the False-Alarm-Rate (FAR) is also evaluated for each candidate. FAR quantifies how often noise can produce a trigger with a ranking statistic $\ln \mathcal{L}$ greater or equal to the ranking statistic $\ln \mathcal{L}^*$ of the trigger under consideration, marginalized over all the data analyzed [11]. Mathematically, we have

$$\text{FAR} = \frac{N \times \text{FAP}}{T}, \quad (4.6)$$

where N is the total number of observed candidates, and T is the duration of the data being analyzed, A ranked list of candidates according to the FAPs estimated is then output from the search pipeline for follow-up analysis like Bayesian parameter estimation [305].

¹ A similar quantity known as the False-Alarm-Rate (FAR) can also be evaluated. FAR tells us how often noise can produce a trigger with a ranking statistic larger or equal to that of the trigger under consideration.

4.2 An overview of the GstLAL search pipeline

In this section, I will give an overview of the GstLAL search pipeline, outlining the key procedures.

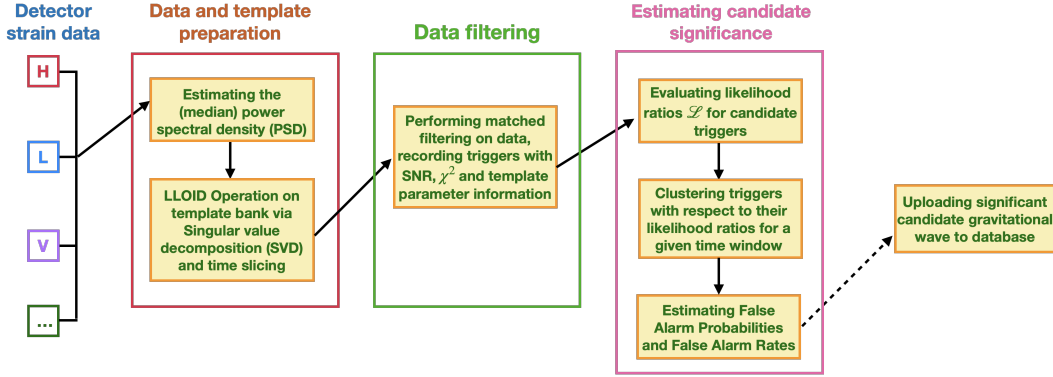


Figure 4.2.1: A simple flowchart outlining the key components of the GstLAL pipeline. Figure reproduced with reference to [250]. Note that the first two boxes representing “data and template preparation” and “data filtering” are done individually for each participating detector.

One key difference that sets GstLAL aside from the other search pipelines like pyCBC is that GstLAL performs matched filtering in the time domain instead of the frequency domain. As we have seen previously, the complex matched-filtering output is given by equation 4.2. GstLAL performs matched-filtering with “real” templates only in the time domain, and hence, for each matched-filtering, GstLAL will only compute the real part x_i of the complex SNR timeseries, corresponding to a single coalescence phase. To obtain the complete complex SNR time series, for each complex template $h_i^c(t)$ (i runs from 0 to $N - 1$, where N is the total number of templates in the template bank, and c means that the template is complex in general), GstLAL will perform two matched-filtering processes with two real sub-templates defined by

$$\begin{cases} h_{2i}(t) = \Re [h_i^c(t)] \\ h_{2i+1}(t) = \Im [h_i^c(t)] \end{cases}, \quad (4.7)$$

where the second waveform differs from the first waveform by a $\frac{\pi}{4}$ coalescence phase shift with respect to Equation 4.1.

Estimating the noise power spectral density (PSD)

As a first step, GstLAL estimates the (one-sided) noise power spectral density (PSD) $S_n(f)$ with the given data from the various detectors. Note that PSD is the square

of the Amplitude Spectral Density (ASD). The PSDs are essential for whitening the templates and data for matched-filtering at a later stage (See also Equation 4.2). Theoretically, if the noise in the detectors is stationary and Gaussian, the one-sided noise PSD is defined [250] as

$$\langle \tilde{n}(f) \tilde{n}^*(f') \rangle = \frac{1}{2} S_n(f) \delta(f - f'), f > 0 \quad (4.8)$$

where the left side is the average of an ensemble of realizations of the noise in the detector $n(t)$. In reality, we can only estimate detector noise from data collected from the detectors, which generally consists of both signals and data. Additionally, detector data is neither stationary nor Gaussian due to glitches. A detailed explanation of how GstLAL estimates the PSDs can be found in Section II A of [250]. As an example, Figure 4.2.2 shows the median power spectral density (PSD) of data collected from the Hanford (H1) and Livingston (L1) detector between UTC time 2023-05-24 15:00:00 and 2023-06-06 14:30:34, estimated using the GstLAL pipeline. The binary-neutron-star horizon distances ² for each detector are also estimated and shown in the legend of the same figure.

Template bank splitting and decomposition

Since we do not know a priori where in the parameter space the gravitational waves we are searching for lives (we refer to this parameter space as the “signal subspace” hereafter), we will need to prepare a large number of gravitational waveforms, collectively stored in a “template bank”, to cover the possible search parameter space where a gravitational wave may live in for a general search. Templates in a standard template bank are not distributed uniformly but are chosen to minimize the size of the template bank. Insufficient templates to cover the parameter space may lead to potential loss of gravitational-wave signals due to excessive loss of signal-to-noise ratio (SNR), causing them not to be identified as potential triggers in the first place in a matched filtering search pipeline. Meanwhile, an over-population of templates leads to high trials factors and thus accumulates excessive “noise” background, making it more difficult to distinguish actual gravitational-wave signals from the collected background ³. Therefore, a standard template bank is generated in a way

² The binary-neutron-star horizon distance is defined as the furthest distance to which we will be able to detect gravitational waves from a $1.4M_\odot - 1.4M_\odot$ binary-neutron-star merger with a threshold signal-to-noise ratio of 8.

³ The trials factor problem (and hence the large noise background) is an obstacle we encountered when looking for sub-threshold (strongly-lensed) gravitational-wave signals, which typically have lower amplitudes and lower ranking statistics. In later chapters, we describe a method to strategically narrow the search parameter space and reduce the noise background to uprank possible sub-threshold gravitational waves.




Figure 4.2.2: The median power spectral density (PSD) of data collected from the Hanford (H1, in red) and Livingston (L1, in blue) detector between UTC time 2023-05-24 15:00:00 and 2023-06-06 14:30:34, estimated using the GstLAL pipeline. The binary-neutron-star horizon distances for the two detectors are also estimated and shown in the legend. Note that GstLAL downsamples the data being analyzed, and hence the PSD shown above ends at 10^3 Hz. However, the PSD extends beyond 10^3 Hz in reality.

that satisfies a minimal match criterion (between any signals within the covered signal subspace and the templates) to balance between identifying signals with minimal loss of SNR and the over-accumulation of the noise background. Methodologies and techniques for creating template banks for searching gravitational waves are out of scope for this thesis, but readers are referred to [80, 96, 310, 359]. Figure 4.2.3 shows the template bank used by the GstLAL search pipeline in LIGO-Virgo-KAGRA's third observing run O3, plotted on the component masses $m_1 - m_2$ space. The entire bank, covering component masses ranging from $1M_\odot$ to $400M_\odot$, is composed of five smaller banks, namely the binary-black-hole (BBH) bank (in blue), binary-neutron-star (BNS) bank (in orange), neutron-star-black-hole (NSBH) bank (in pink), intermediate-mass-black-hole (IMBH) bank (in green) and other BBH bank (in red). They are used to target different regions in the signal subspace. The bank has a total of 1,412,263 templates.

The large number of templates in the template bank still poses the problem of high computational cost induced by running matched filtering with all the templates.




Figure 4.2.3: The template bank used by the GstLAL search pipeline in O3, shown in the component masses $m_1 - m_2$ space. The full bank is composed of five smaller banks, including the binary-black-hole (BBH) bank (in blue), binary-neutron-star (BNS) bank (in orange), neutron-star-black-hole (NSBH) bank (in pink), intermediate-mass-black-hole (IMBH) bank (in green) and other BBH bank (in red). These different banks are used to target different regions in signal subspace.

GstLAL addresses this problem via the Low Latency Online Inspiral Detection (LLOID) method [87]. First, GstLAL splits the full template bank into smaller “split-banks”, each carrying $2N_T$ real templates⁴, according to the templates’ chirp mass

$$\mathcal{M}_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (4.9)$$

and effective spin parameter

$$\chi_{\text{eff}} = \frac{m_1 s_{1z} + m_2 s_{2z}}{m_1 + m_2}, \quad (4.10)$$

where m_1 , m_2 are the component masses, and s_{1z} and s_{2z} are the components’ dimensionless spins in the z -direction. The neighbouring split banks in the \mathcal{M}_c

⁴ The factor of 2 is the result of GstLAL using only real templates to compute the complex SNR timeseries; N_T is an adjustable number which is generally set to be $O(100)$.

space are partially overlapping for intermediate use when adopting the LLOID method, but will be removed later so there are no repeated templates at the end. Figure 4.2.4 shows an example of two split banks (bank 1000 and 1019) generated from the full template bank (Figure 4.2.3) in the chirp mass \mathcal{M}_c - effective spin χ_{eff} space.




Figure 4.2.4: An example of two split banks (bank 1000 and 1019) generated from the full template bank, shown in the chirp mass \mathcal{M}_c - effective spin χ_{eff} space. Note that the partially overlapping region exists only for intermediate use when adopting the LLOID method, but will be removed afterward so that there are no repeated templates.

GstLAL then sub-divides each sub-bank into smaller “time slices”, so each waveform is further “chopped” into smaller regions in time with respect to the waveform frequencies at each time. The sampling frequency for each time slice will be adjusted according to the corresponding Nyquist frequency. Hence, we can use a lower sampling frequency for the low-frequency time slices, reducing the number of sampling points and, thus, the computation cost when performing matched filtering. Figure 4.2.5 gives an example of the timeslices for a template waveform with chirp mass $\mathcal{M}_c = 2.61 M_\odot$ and z-component spins $s_{1z} = -0.73$ and $s_{2z} = -1.0$.




Figure 4.2.5: An example of the time slices for a template waveform with chirp mass $\mathcal{M}_c = 2.61M_\odot$ and z-component spins $s_{1z} = -0.73$ and $s_{2z} = -1.0$. Note that the sampling frequency increases with time as the Nyquist frequency increases.

The final step in the LLOID method is to perform a singular-value decomposition (SVD) on each time slice for each sub-bank. SVD reduces the physical templates into unphysical templates (i.e. they look nothing like actual gravitational waveforms) as basis. One can reconstruct the original waveform with the corresponding reconstruction coefficients (also known as the mixing matrix elements). GstLAL truncates the SVD output so that it only keeps the most important basis to further reduce computational cost with minimal effect on the waveform reconstruction accuracy (set by the “tolerance”). In the end, GstLAL will perform matched filtering with the SVD basis, and reconstruct the SNR timeseries for each template waveform. The process is shown schematically in figure 4.2.6. The mathematical derivations and computation of the LLOID method can be found in [87].

Matched-filtering and the signal consistency test

When performing matched filtering, GstLAL records instances when the SNR exceeds some pre-determined threshold as candidate events (also known as “triggers”).

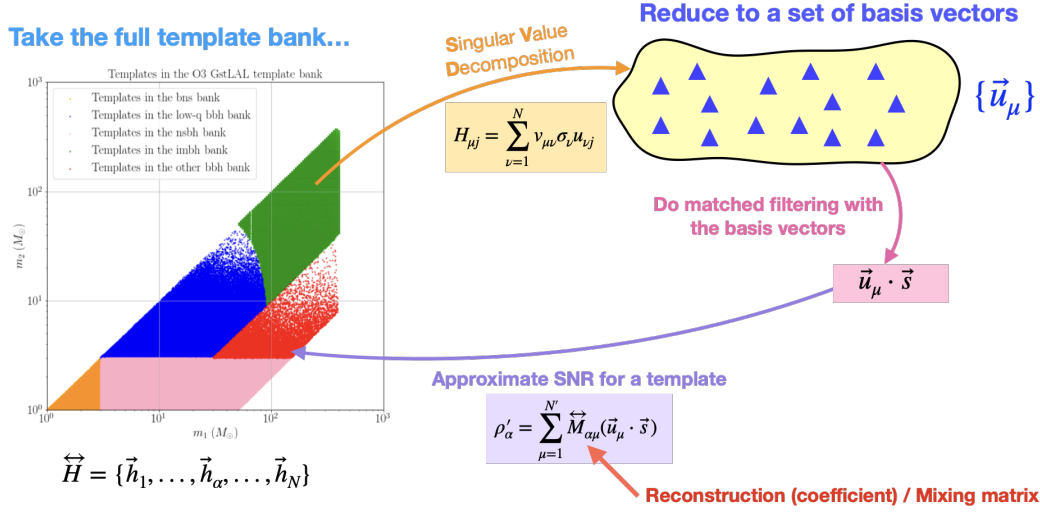


Figure 4.2.6: A schematic diagram of the singular-value-decomposition (SVD) operation on the template bank.

At the time of writing, GstLAL sets the SNR threshold for a single detector trigger at 4.

Suppose noise in the data is truly Gaussian and stationary. In that case, the matched-filtered SNR will be an optimal ranking statistic to determine how likely a trigger is an actual gravitational-wave event compared to noise. However, this is not the case for real data. Noise fluctuations and non-Gaussianities (e.g. glitches) in the data, can also produce high SNR triggers in the pipeline. To address this problem, GstLAL assigns a second detection statistic known as the “auto-correlation based signal consistency test value” ξ^2 to each trigger. GstLAL precomputes the expected complex auto-correlation timeseries for each template waveform (i.e. the expected SNR timeseries when a signal is matched-filtered against a template waveform having the same parameters) as [250]

$$R_j(t) = \int \frac{|\tilde{h}_{2j}(f)|^2 + |\tilde{h}_{2j+1}(f)|^2}{S_n(f)} e^{2\pi i f t} df, \quad (4.11)$$

where j runs from 0 to $N_T - 1$, corresponding to a total of N_T physical templates. h_{2j} and h_{2j+1} are the real and imaginary parts of the template waveform defined previously. We can compare these to the actual obtained complex SNR timeseries for a trigger

$$z_j(t) = x_{2j}(t) + i x_{2j+1}(t), \quad (4.12)$$

Conventionally, we set $t = 0$ to be the peak time, and the autocorrelation is normalized for each template such that $R(0) = 1$. If we assume that the actual SNR




Figure 4.2.7: The auto-correlation timeseries (blue striped line) and SNR timeseries (black solid line) evaluated for GW230529 for the LIGO Livingston (L1) detector. In the above figure, the trigger’s coalescence time is set to $t = 0$.

timeseries of a signal that is identical to a template waveform can be represented by the corresponding autocorrelation timeseries up to a scaling factor $z_j(0)$ ⁵, then we can compute ξ^2 as

$$\xi^2 = \frac{\int_{-\delta t}^{\delta t} |z(t) - z(0)R(t)|^2 dt}{\int_{-\delta t}^{\delta t} 2 - 2|R(t)|^2 dt}. \quad (4.13)$$

The numerator is effectively the sum of the squared difference between the expected (auto-correlated) normalized SNR timeseries $R_j(t)$ (scaled by the peak SNR of the actual SNR timeseries) and the actual SNR timeseries $z_j(t)$, integrated over a predetermined time window δt around the peak time. Details about how δt is set can be found in [250]. The denominator of ξ^2 is merely a normalization constant, such that a perfect match between the autocorrelation SNR timeseries and actual SNR timeseries will give $\xi^2 = 1$. In principle, the larger ξ^2 is, the more different the trigger’s SNR timeseries is as compared to the expected SNR timeseries, and

⁵ Here, $t = 0$ is again chosen to be the peak time of the actual SNR timeseries such that $\max(z_j(t)) = z_j(0)$.

hence, the less likely it is a gravitational wave that is described by the template waveform ⁶. Figure 4.2.7 shows an example of the autocorrelation timeseries (blue) and actual timeseries (black) for the event GW230529 for the LIGO Livingston (L1) detector. In the Appendix Section 4.A, I also provide a detailed derivation for the normalization constant as reference.

Collecting signal candidates and noise candidates

After filtering, the GstLAL pipeline now has a collection of triggers, each of which contains (1) the trigger time t , (2) the template parameters $\vec{\theta}_i$ as a label, (3) the matched-filtered SNR ρ , and (4) the auto-correlation-based signal consistency test value ξ^2 as trigger information. It is important to note that at a given time instant, multiple triggers can be formed with different templates in the template bank (i.e. triggers at a given timestamp do not have unique template parameters, hence SNR and ξ^2).

Starting with the collection of triggers, GstLAL looks for triggers that are found in time coincidence (within a certain time window that accounts for (1) the maximum travel time of gravitational waves between detectors, and (2) statistical fluctuations in the measured trigger time due to detector noise) between participating detectors with the same template parameters. These triggers are grouped and labelled as coincident (signal) triggers [250]. Coincident triggers are further followed up by the pipeline as signal triggers. In contrast, single detector triggers (non-coincident triggers) are treated as noise (background) triggers and are used to construct the noise background. Starting from LVK's second observing run O2, GstLAL also allows single detector triggers to be treated as signal candidates if they pass a preliminary log-likelihood ratio threshold. This will be further explained in later subsections (Or, see Section D1 of [309]).

Key terms in the Likelihood Ratio Statistic

A ranking statistic needs to be assigned to each of the triggers to rank candidate events from least likely to be a real signal, and most likely to be a real signal. GstLAL uses

⁶ While ξ^2 can help the pipeline to distinguish triggers corresponding to real gravitational-wave signals from triggers corresponding to glitches, one may also potentially lose exotic gravitational-wave signals that do not look like the template waveforms. For example, suppose there exists an alternative-gravity theory (Non-GR) gravitational-wave signal in the data that differs significantly from usual general-relativity gravitational template waveforms; the pipeline will assign a high ξ^2 to the trigger (if the signal is registered as a trigger in the first place), and will potentially affect its ranking statistics at a later stage, making it a sub-threshold candidate at the end. Another example would be microlensed gravitational waveforms (See Part II of the thesis).




Figure 4.2.8: Triggers found in time coincidence (within a certain time window that accounts for (1) the maximum travel time of gravitational waves between detectors, and (2) statistical fluctuations in the measured trigger time due to detector noise) between participating detectors with the same template parameters are grouped and labelled as coincident (signal) triggers by GstLAL. Coincident triggers are further followed up by the pipeline as signal triggers. In contrast, single detector triggers (non-coincident triggers) are treated as noise (background) triggers by GstLAL, and are used to construct the noise background. Starting from LVK’s second observing run O2, GstLAL also allows single detector triggers to be treated as signal candidates if they pass a preliminary log-likelihood ratio threshold. This will be further explained in later subsections.

the likelihood-ratio ranking statistic [85, 166, 250, 309] to rank candidate events. For coincident triggers, the likelihood ratio $\mathcal{L}_{\text{coinc}}$ is defined as

$$\mathcal{L} = \frac{P(\vec{D}_H, \vec{O}, \vec{\rho}, \vec{\xi}^2, [\vec{\Delta}t, \vec{\Delta}\phi] | \vec{\theta}, \text{signal})}{P(\vec{D}_H, \vec{O}, \vec{\rho}, \vec{\xi}^2, [\vec{\Delta}t, \vec{\Delta}\phi] | \vec{\theta}, \text{noise})} \cdot \frac{P(\vec{\theta} | \text{signal})}{P(\vec{\theta} | \text{noise})}, \quad (4.14)$$

where (1) \vec{D}_H are the horizon distances of the detectors (a measure of the detector sensitivity), (2) \vec{O} is the set of detectors that registered the coincident trigger, (3) $\vec{\rho}$ are the recorded detector SNRs for the trigger at each participating detectors, (4) $\vec{\xi}^2$ are the auto-correlation-based consistency test values evaluated for each of the participating detectors, (5) $[\vec{\Delta}t, \vec{\Delta}\phi]$ are the sets of differences in end times Δt and coalescence phases $\Delta\phi$ between pairs of detectors within the set of participating detectors⁷, and (6) $\vec{\theta}$ is the template parameters associated with the trigger.

The likelihood ratio $\mathcal{L}_{\text{coinc}}$, in essence, is the ratio of the probability of getting the sets of detection statistics for a given coincident triggers under the hypothesis that the trigger is a signal (signal hypothesis), to that under the hypothesis that trigger

⁷ For instance, if a coincident trigger is found by three detectors, H1, L1, and V1, then $\vec{\Delta}t$ will have three components, Δt_{H1L1} , Δt_{L1V1} and Δt_{H1V1} , corresponding to the difference in end times between the H1-L1 detector pair, L1-V1 detector pair and H1-V1 detector pair.

originates from noise (noise hypothesis). The higher the likelihood ratio, the more likely the trigger is an actual gravitational-wave signal.

Starting from O2, GstLAL also computes an alternative likelihood ratio for single detector triggers [309] as

$$\mathcal{L}_{\text{single}} = \frac{P(\vec{D}_H, O, \rho, \xi^2 | \vec{\theta}, \text{signal})}{P(\vec{D}_H, O, \rho, \xi^2 | \vec{\theta}, \text{noise})} \cdot \frac{P(\vec{\theta} | \text{signal})}{P(\vec{\theta} | \text{noise})}, \quad (4.15)$$

that is mostly the same as the likelihood ratio for coincident triggers, except that (1) the SNRs and ξ^2 vectors now become only the values for the detector that registers the trigger, (2) the difference in end times and coalescence phases $[\vec{\Delta}t, \vec{\Delta}\phi]$ terms are dropped, (3) the set of participating detectors \vec{O} now reduces to just the detector that registers the trigger. An additional penalty of 12 is manually subtracted from the **log** single detector likelihood ratio [153] to account for single detector triggers being more susceptible to uncertainty as they are not found in time coincidence in multiple detectors. For single detector triggers that have a log-likelihood ratio $\ln \mathcal{L}_{\text{single}} > 2$, GstLAL will consider them as signal triggers for further follow-up.

As explained in [85, 250], the likelihood ratio can be factorized into smaller terms. Details of the factorization and the calculation for each of the sub-terms can be found in [85, 166, 250, 309]. Here, I will discuss several key terms in preparation for the discussion in later Chapters.

The $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ term

$P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ is the probability of registering a (coincident) noise trigger with SNRs $\vec{\rho}$ and ξ^2 . To compute this probability, GstLAL first populates histograms in the ρ - ξ^2 space with collected background triggers for each detector. The histograms are then normalized⁸ and smoothed with a Gaussian smoothing kernel. Figure 4.2.9 shows an example histogram for noise triggers collected in the Hanford (H1) detector for one split bank between GPS time 1370097052 and 1370356252 (3 days of data). These histograms do not track time dependence in both the online and offline

⁸ The ρ - ξ^2 histogram contains 2D bins of widths $\Delta\rho_i$ and $\Delta\xi_i^2$ for the i^{th} bin. To normalize the histogram, one “integrates” the histogram via

$$C = \sum_i n_i \Delta\rho_i \Delta\xi_i^2, \quad (4.16)$$

where i runs from 0 to N , the total number of bins, and n_i represents the counts in the i^{th} bin, to obtain the normalization constant C . Then, the histogram bin counts are divided by the normalization constant. This ensures that the normalized histogram, when integrated, gives 1 as any usual probability density.




Figure 4.2.9: The ρ - ξ^2 histogram for noise triggers collected in the Hanford (H1) detector for one split bank between GPS time 1370097052 and 1370356252 (3 days of data). The green circle marks an example signal trigger with H1 SNR $\rho_{\text{H1}} = 10$ and $\xi_{\text{H1}}^2 = 1$. Since it is (marginally) away from the background noise triggers distribution, GstLAL will evaluate a low $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ for this trigger, in turn increasing the log-likelihood ratio and hence the significance of this trigger.

operation, but instead, they are merely cumulative histograms. This will potentially affect the sensitivity of the search because of durations of noisy data, which can occur when detectors are affected by seismic motion and extreme weather, can pollute the ρ - ξ^2 histograms. Should a real signal trigger be found later during relatively “cleaner” data, the trigger will still be assigned a $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ probability using the polluted histogram, reducing its likelihood ratio and significance. This may potentially result in the pipeline not “finding”⁹ the trigger, in turn reducing the search sensitivity of the pipeline. In Chapter 5, we explore and discuss a potential solution to tackle this problem.

The $P(\vec{\Delta}t | \vec{\rho}, \vec{D}_H, \text{signal}) \times P(\vec{\Delta}\phi | \vec{\Delta}t, \vec{\rho}, \vec{D}_H, \text{signal})$ term

The $P(\vec{\Delta}t | \vec{\rho}, \vec{D}_H, \text{signal}) \times P(\vec{\Delta}\phi | \vec{\Delta}t, \vec{\rho}, \vec{D}_H, \text{signal})$ term is added to the likelihood ratio calculations since LVK’s second observing run O2. This term accounts for

⁹ A signal trigger is “found” if its ranking statistic/significance passes a pre-determined threshold. For the GstLAL pipeline it thresholds on the False-Alarm-Rate (FAR) of a trigger (To be explained in later subsections).

signal consistency across the detector network. The arrival times and coalescence




Figure 4.2.10: (Left) At different sky locations, a gravitational wave may reach one detector before reaching the other. The effective strength of the signal may also differ. Depending on the ratio of the signal strengths (effective distance ratios), the expected time and phase delay distribution for a pair of detectors has a non-trivial probability distribution, given that they are seeing the same signal. (Right) A cartoon illustrating an example of the Δt , $\Delta \phi$ probability distribution for a trigger seen in both the Hanford (H1) and Livingston (L1) detector with an identical SNR of 8. Assuming both detectors have similar sensitivity, it is more likely for the signal to be detected at the same time with almost no difference in the coalescence phase. Note that this is just a cartoon illustration. See Figure 7 of [309] for a realistic example.

phases of signal triggers at each detector, assuming sources of gravitational waves are distributed isotropically across the sky, is uniform. However, the relative time delays and coalescence phase of the same signal arriving at different detector make up a non-trivial correlated probability density function, depending on the ratio of the signal strengths (effective distance ratios). It is, therefore, natural to precompute the probability distributions for signals and use them to check the signal consistency for signal triggers across detectors. [183, 309] give detailed discussion and explanation on how to precompute the aforementioned term.

It is, however, helpful to note that the current precomputed probability distributions are marginalized over all sky and all possible inclinations and polarizations. This marginalization prevents GstLAL from performing any sky-location-based targeted searches. For instance, one can search for gravitational-wave counterparts to a Gamma-Ray-Burst (GRB) detection. The targeted search can be made more sensitive if we can focus on regions in the sky that overlap with the localization of the GRB source. In later chapters (Chapter 13), we will also see that sky-location-based targeted searches will be useful for uncovering possible strongly-lensed gravitational waves coming from the same source since they can be treated as having the same sky location essentially.

As part of ongoing work, we are investigating the possibility of rewriting the log-likelihood ratio so that the sky localization of the source (i.e. the right ascension α and declination δ) can be included as part of the calculations, in replacement of the $\vec{\Delta}t, \vec{\Delta}\phi$ term. This will become exceptionally useful as more detectors (e.g. Virgo, KAGRA and LIGO India) join the detector network. To do so, we are now working on implementing the rapid localization method introduced in [347] into the GstLAL search pipeline as a first step. In short, [347] creates posterior skymaps for the sources of gravitational-wave signals in a similar fashion as existing gravitational-wave source localization pipelines like BAYESTAR [319], but it **maximizes**¹⁰ the posterior skymap over the source distance and the source’s inclination. While sacrificing a certain amount of accuracy in the skymap, this maximization increases the computation speed as part of the terms in the computation becomes data-independent and can be factorized out and precomputed. According to [347], their method reduces the computation time of a skymap to $\mathcal{O}(0.1s)$ as compared to $\mathcal{O}(1s)$ for BAYESTAR [319]. We aim to implement this method into GstLAL so that the pipeline will evaluate the source localization for signal triggers as an additional statistic for computing the likelihood ratio. However, work is still ongoing to (1) determine what parameters/terms should be used to represent the source localization of a trigger and (2) develop a (semi-analytic) signal and noise model for the sky localization parameters we are introducing to the likelihood ratio calculation, and (3) further improve the accuracy of the source localization method introduced in [347]. We also note that the computational time/cost for computing skymaps for every single trigger coming out of the pipeline can still be large, given the large number of triggers (See also “Fast Path Cut” and the subsection “Clustering”). We will need to develop a more effective way to pre-select candidates that are more likely to be gravitational-wave signals so that we can drastically reduce the number of triggers the pipeline needs to compute source localization for. This pivotal change in the likelihood ratio calculation will allow GstLAL to perform sky-location based targeted searches in the future.

The $P(\vec{\theta}|\text{signal})$ term

The $P(\vec{\theta}|\text{signal})$ term is the numerator in the second part of the likelihood ratio definition. This term, formally known as the **population model term**, corresponds to the probability that a template waveform with parameters $\vec{\theta}$ is recovered as the

¹⁰ For reference, BAYESTAR [319] marginalizes the posterior skymaps over source distance and inclination.

best-match template, given that the trigger is a signal. Before LVK's third observing run O3, the population model term was neglected and is set to be 1 for every template in the template bank [166, 309]. Effectively, this assumes that astrophysical sources are distributed identically to how the templates are distributed in the template bank. This is incorrect, mainly due to the uneven density of templates in the template bank.

[166] provides a detailed discussion on properly incorporating prior information about the signal distribution in parameter space (population model) into the population model term. Here, we give a summary for completeness. We start by assuming the following:

- (1) The template bank is populated with a sufficiently high density such that any gravitational-wave signal $\vec{s}(\vec{\theta})$ labelled by its source parameters $\vec{\theta}$ can be described exactly by a normalized template \vec{t}_j up to a proportionality constant ρ that describes the strength of the signal (i.e. the signal-to-noise ratio SNR). Mathematically, we have

$$\vec{s}(\vec{\theta}) \sim \rho \vec{t}_j. \quad (4.17)$$

- (2) Noise in the data is stationary and Gaussian, and
- (3) Data \vec{d} that contains both noise \vec{n} and a gravitational-wave signal \vec{s} can be described by the multiple of a normalized waveform template \vec{t}_k , i.e.

$$\vec{d} \sim \rho_{\text{obs},k} \vec{t}_k. \quad (4.18)$$

Suppose there is a gravitational-wave signal $\vec{s} = \rho \vec{t}_j$ originating from a predefined population in the data, then the probability for the signal to be picked up by a template \vec{t}_k is given by

$$P(\vec{t}_k | s, \rho) = \sum_j \{P(\text{candidate in } \vec{t}_k | s = \rho \vec{t}_j) \times P(\vec{t}_j | \text{signal from population})\}, \quad (4.19)$$

where j runs from 0 to N , with N being the total number of templates in the template bank. The first term in the summation, $P(\text{candidate in } \vec{t}_k | s = \rho \vec{t}_j)$, gives the probability that the observed signal described by \vec{t}_j is recovered as a trigger by template \vec{t}_k . This term can be precomputed as the **template overlap** for each template in the template bank. The second term, $P(\vec{t}_j | \text{signal from population})$,

gives the probability that a signal originating from the pre-determined population can be described by a template \vec{t}_j . One can incorporate a population model into the likelihood ratio calculation in this term. Details of the exact computation can be found in [166]. Nevertheless, it is useful to note that in later chapters (Chapter 11), we will show how one can incorporate a targeted population model to improve the search sensitivity for possible sub-threshold strongly lensed gravitational-wave signals.

Fast path cut

While not adequately documented in any of the published articles so far, GstLAL actually employs a “fast path cut” function to “discard” triggers that analysts deemed unlikely to be an actual signal. The logic of the function goes as follows: Given a coincident trigger, if the network SNR ¹¹ of the trigger is < 7 , the trigger will be automatically assigned a log-likelihood ratio of negative infinity. This is essentially “discarding” the trigger. The reason for employing the “fast-path cut” function is to reduce the number of triggers for follow-up preliminarily. However, this leads to two potential issues: (1) The $\rho_{\text{network}} \geq 7$ threshold is arbitrarily set. There has been no in-depth studies done a priori before this threshold is set. Also, as we have mentioned earlier, SNR is not an optimal ranking statistic to determine how likely a trigger is a real gravitational-wave signal, given that noise is not stationary and Gaussian. The “fast-path cut” function can potentially reduce the search sensitivity by unknowingly removing signal triggers. (2) The “fast path cut” function, together with trigger clustering (to be discussed in the next subsection) will lead to permanent loss of trigger information because the pipeline will ultimately remove triggers ¹² with low log-likelihood ratios. This prevents analysts from utilizing filtered products (the complete trigger list) from a given search for alternative studies. For instance, one may want to rerank ¹³ triggers from a search using a different population model. Another example would be to rerank candidates from multiple searches across different time ranges of data with a combined background. These would not be possible since some triggers have been clustered and discarded from the complete list of triggers, potentially leading to inaccurate results from the re-analyses.

¹¹ The network SNR of a trigger is the SNRs recorded in all participating detectors added in quadrature, i.e. $\rho_{\text{network}} = \sqrt{\rho_i \rho_i}$.

¹² Here, removing means deleting files and information related to a trigger.

¹³ A rerank involves a re-assignment of log-likelihood ratios to the triggers.

Clustering

As mentioned before, different templates in the template bank can form multiple (coincident) signal triggers at a given time. [250] also mentioned that it is probable for GstLAL to register at least one trigger per every one-second interval per template per detector. To reduce the size of the final data products (i.e., the trigger list and associated files), GstLAL keeps only triggers with the highest likelihood ratio within a certain time window and permanently discards the other triggers. This process is known as **clustering**. For the online analysis, clustering is done for each split bank within a ± 1 -second window. For the offline analysis, clustering is done within a ± 4 -second window across the full template bank. As explained in the previous subsection, the clustering process prevents analysts from using the data products in a search for alternative studies.

As part of ongoing work, we are investigating an alternative, better method/logic to replace the current “fast path cut” function, such that we can still retain the complete list of triggers but limit the follow-up analysis to only those that are (semi-analytically) more likely to be real gravitational-wave signals. The aim is to trim down the list of triggers for follow-up without discarding any potential ones that may be real signals to a point where the size of the final data products is manageable, such that the clustering process can be removed and information of the complete list of triggers can be kept. In the future, this will allow users to properly do reranks and other studies using data products from GstLAL searches, without the need to do re-filtering (i.e. a new search).

False Alarm Rates and False Alarm Probabilities

As a final step, GstLAL evaluates, for each trigger, a False-Alarm-Probability (FAP) defined as

$$\text{FAP} = P(\ln \mathcal{L} > \ln \mathcal{L}^* | \text{noise}) = \int_{\ln \mathcal{L}^*}^{\infty} P(\ln \mathcal{L} | \text{noise}) d \ln \mathcal{L}. \quad (4.20)$$

Effectively, the FAP is the probability that noise can produce a trigger with log-likelihood ratio $\ln \mathcal{L}$ larger or equal to the log-likelihood ratio $\ln \mathcal{L}^*$ of the trigger under consideration [250]¹⁴. We can, from the FAP, evaluate a similar statistic known as the False-Alarm-Rate (FAR) defined by

$$\text{FAR} = \frac{\text{FAP} \times N}{T}, \quad (4.21)$$

¹⁴ The False-Alarm-Probability is also known more generally as the p-value.

where N is the total number of candidates, and T is the total duration of data being analyzed. The FAR tells us how often noise can produce a trigger with log-likelihood ratio $\ln \mathcal{L}$ larger or equal to the log likelihood ratio $\ln \mathcal{L}^*$ of the trigger under consideration [250]. The lower the FAR, the more significant a trigger is.

Ultimately, GstLAL outputs a list of triggers ranked by their FARs. Candidates that pass a pre-determined FAR threshold are passed on for follow-up analyses, including Bayesian parameter estimation.

More information about the GstLAL pipeline can be found in [86, 250, 309].

APPENDIX

4.A Deriving the normalization constant for GstLAL's auto-correlation based signal consistency test value ξ^2

Recall that the auto-correlation signal-consistency test value without normalization is given by

$$\tilde{\xi}^2(t) = |z_j(t) - z_j(0)R_j(t)|^2, \quad (4.22)$$

where $z_j(t)$ is the complex SNR time series

$$z_j(t) = 2 \int_{-\infty}^{+\infty} \frac{[h_{2j}(f) + ih_{2j+1}(f)]^* d(f)}{S_n(|f|)} e^{2\pi i f t} df, \quad (4.23)$$

with h_{2j} and h_{2j+1} being the real and imaginary part of the template waveform, $d(f)$ is the data in the frequency domain, $R_j(t)$ is the normalized autocorrelation function of the template waveform, $S_n(|f|)$ is the noise power spectral density.

The expectation value of $\tilde{\xi}^2$ is given by

$$\langle \xi^2 \rangle = \langle |z_j(t) - z_j(0)R_j(t)|^2 \rangle \quad (4.24)$$

$$= \langle |z_j(t)|^2 + |z_j(0)|^2 |R_j(t)|^2 - z_j(t)z_j^*(0)R_j^*(t) - z_j^*(t)z_j(0)R_j(t) \rangle \quad (4.25)$$

$$= \langle |z_j(t)|^2 \rangle + \langle |z_j(0)|^2 |R_j(t)|^2 \rangle - \langle z_j(t)z_j^*(0)R_j^*(t) \rangle - \langle z_j^*(t)z_j(0)R_j(t) \rangle. \quad (4.26)$$

This is the expectation value when the data is purely Gaussian noise, i.e. $d(f) = n(f)$. In the frequency domain, the auto-correlation function $R_j(t)$ can be written as

$$R_j(t) = \int_{-\infty}^{\infty} \frac{|h_{2j}(f)|^2 + |h_{2j+1}(f)|^2}{S_n(|f|)} e^{2\pi i f t} df. \quad (4.27)$$

With this, part of the second term in the expanded version of $\tilde{\xi}^2$ can be evaluated as

$$\langle |z_j(0)|^2 \rangle = \langle 4 \underbrace{\int_{-\infty}^{\infty} \frac{[h_{2j}(f_1) + ih_{2j+1}(f_1)]n^*(f_1)}{S_n^*(|f_1|)} df_1}_{z_j^*(0)} \underbrace{\int_{-\infty}^{\infty} \frac{[h_{2j}(f_2) + ih_{2j+1}(f_2)]^* n(f_2)}{S_n(|f_2|)} df_2}_{z_j(0)} \rangle \quad (4.28)$$

$$= 4 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \langle n^*(f_1) n(f_2) \rangle \left\langle \frac{[h_{2j}(f_1) + ih_{2j+1}(f_1)][h_{2j}^*(f_2) - ih_{2j+1}^*(f_2)]}{S_n^*(|f_1|) S_n(|f_2|)} \right\rangle df_1 df_2 \quad (4.29)$$

$$= 2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} S_n(|f_2|) \delta(f_2 - f_1) \left\langle \frac{[h_{2j}^*(f_1) + ih_{2j+1}^*(f_1)][h_{2j}^*(f_2) - ih_{2j+1}^*(f_2)]}{S_n^*(|f_1|) S_n(|f_2|)} \right\rangle df_1 df_2 \quad (4.30)$$

$$= 2 \int_{-\infty}^{\infty} \left\langle \frac{[h_{2j}^*(f_1) + ih_{2j+1}^*(f_1)][h_{2j}^*(f_1) - ih_{2j+1}^*(f_1)]}{S_n^*(|f_1|)} df_1 \right\rangle \quad (4.31)$$

$$= 2 \underbrace{\int_{-\infty}^{\infty} \frac{|h_{2j}(f_1)|^2 + |h_{2j+1}(f_1)|^2}{S_n^*(|f_1|)} df_1}_{R_j^*(0)=1} \quad (4.32)$$

$$= 2, \quad (4.33)$$

where we have made use of the fact that for stationary Gaussian noise,

$$\langle n(f_1) n^*(f_2) \rangle = \frac{1}{2} S_n(f_1) \delta(f_1 - f_2). \quad (4.34)$$

Similarly, the first term can be evaluated as

$$\langle |z_j(t)|^2 \rangle = \langle 4 \underbrace{\int_{-\infty}^{\infty} \frac{[h_{2j}(f_1) + ih_{2j+1}(f_1)]n^*(f_1)}{S_n^*(|f_1|)} e^{-2\pi i f_1 t} df_1}_{z_j^*(t)} \underbrace{\int_{-\infty}^{\infty} \frac{[h_{2j}(f_2) + ih_{2j+1}(f_2)]^* n(f_2)}{S_n(|f_2|)} e^{2\pi i f_2 t} df_2}_{z_j(t)} \rangle \quad (4.35)$$

$$= 4 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \langle n^*(f_1) n(f_2) \rangle \left\langle \frac{[h_{2j}(f_1) + ih_{2j+1}(f_1)][h_{2j}^*(f_2) - ih_{2j+1}^*(f_2)]}{S_n^*(|f_1|) S_n(|f_2|)} e^{2\pi i (f_2 - f_1) t} \right\rangle df_1 df_2 \quad (4.36)$$

$$= 2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} S_n(|f_2|) \delta(f_2 - f_1) \left\langle \frac{[h_{2j}(f_1) + ih_{2j+1}(f_1)][h_{2j}^*(f_2) - ih_{2j+1}^*(f_2)]}{S_n^*(|f_1|) S_n(|f_2|)} e^{2\pi i (f_2 - f_1) t} \right\rangle df_1 df_2 \quad (4.37)$$

$$= 2 \int_{-\infty}^{\infty} \left\langle \frac{[h_{2j}(f_1) + ih_{2j+1}(f_1)][h_{2j}^*(f_1) - ih_{2j+1}^*(f_1)]}{S_n^*(|f_1|)} df_1 \right\rangle \quad (4.38)$$

$$= 2 \underbrace{\int_{-\infty}^{\infty} \frac{|h_{2j}(f_1)|^2 + |h_{2j+1}(f_1)|^2}{S_n^*(|f_1|)} df_1}_{R_j^*(0)=1} \quad (4.39)$$

$$= 2. \quad (4.40)$$

Note that

$$\langle z_j(t) z_j^*(0) \rangle = \underbrace{\left\langle 4 \int_{-\infty}^{\infty} \frac{[h_{2j}(f_1) + i h_{2j+1}(f_1)]^* n(f_1)}{S_n(|f_1|)} e^{2\pi i f_1 t} df_1 \right\rangle}_{z_j(t)} \underbrace{\left\langle \int_{-\infty}^{\infty} \frac{[h_{2j}(f_2) + i h_{2j+1}(f_2)] n^*(f_2)}{S_n^*(|f_2|)} df_2 \right\rangle}_{z_j^*(0)} \quad (4.41)$$

$$= 4 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \langle n(f_1) n^*(f_2) \rangle \left\langle \frac{[h_{2j}^*(f_1) - i h_{2j+1}^*(f_1)] [h_{2j}(f_2) + i h_{2j+1}(f_2)]}{S_n(|f_1|) S_n^*(|f_2|)} e^{2\pi i f_1 t} \right\rangle df_1 df_2 \quad (4.42)$$

$$= 2 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} S_n(|f_1|) \delta(f_1 - f_2) \left\langle \frac{[h_{2j}^*(f_1) - i h_{2j+1}^*(f_1)] [h_{2j}(f_2) + i h_{2j+1}(f_2)]}{S_n(|f_1|) S_n^*(|f_2|)} e^{2\pi i f_1 t} \right\rangle df_1 df_2 \quad (4.43)$$

$$= 2 \int_{-\infty}^{\infty} \left\langle \frac{[h_{2j}^*(f_2) - i h_{2j+1}^*(f_2)] [h_{2j}(f_2) + i h_{2j+1}(f_2)]}{S_n^*(|f_2|)} e^{2\pi i f_2 t} df_2 \right\rangle \quad (4.44)$$

$$= 2 \underbrace{\int_{-\infty}^{\infty} \frac{|h_{2j}(f_2)|^2 + |h_{2j+1}(f_2)|^2}{S_n^*(|f_2|)} e^{2\pi i f_2 t} df_2}_{R_j(t)} \quad (4.45)$$

$$= 2R_j(t), \quad (4.46)$$

and similarly.

$$\langle z_j^*(t) z_j(0) \rangle = 2R_j^*(t). \quad (4.47)$$

Putting all these results together, we get

$$\langle \xi^2 \rangle = \langle |z_j(t)|^2 \rangle + \langle |z_j(0)|^2 |R_j(t)|^2 \rangle - \langle z_j(t) z_j^*(0) R_j^*(t) \rangle - \langle z_j^*(t) z_j(0) R_j(t) \rangle \quad (4.48)$$

$$= 2 + 2|R_j(t)|^2 - 2|R_j(t)|^2 - 2|R_j(t)|^2 \quad (4.49)$$

$$= 2 - 2|R_j(t)|^2 \quad (4.50)$$

which is the normalization term for ξ^2 .

Chapter 5

IMPROVING THE RANKING STATISTIC ASSIGNMENT OF THE GSTLAL SEARCH PIPELINE WITH THE INCLUSION OF IDQ DATA-QUALITY INFORMATION

Note: This chapter will be turned into a short author publication:

“Improving the ranking statistic assignment of the GstLAL search pipeline with the inclusion of iDQ data-quality information”, In preparation.

5.1 Introduction and Overview

One of the GstLAL search pipeline’s main development goals is to improve its search sensitivity towards gravitational waves from compact binary coalescences. As we have briefly mentioned in the previous chapter, the lack of time dependence of the ρ - ξ^2 histograms used by the GstLAL pipeline to assign ranking statistics can potentially cause GstLAL to miss gravitational-wave signals, leading to a loss in search sensitivity.

In this chapter, I describe a method that utilizes statistical data quality from the iDQ pipeline [149, 176] information as a workaround to the stated problem above. This chapter is arranged as follows: In Section 5.2, I will briefly recap the computation of the $P(\rho, \xi^2 | \text{noise})$ portion for the calculation of the likelihood ratio for the GstLAL pipeline and give a short introduction to iDQ, a machine-learning-based data quality pipeline. In Section 5.3, I will outline the proposed method to incorporate information from the iDQ pipeline into GstLAL’s likelihood ratio calculation. Finally, in Section 5.4, I will briefly discuss the expected results with the successful implementation of the proposed method.

5.2 Background

In this section, I will briefly recap how GstLAL computes the likelihood ratio for triggers, explicitly focusing on the $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ term. I will briefly introduce the iDQ data quality pipeline [149, 176].

GstLAL's likelihood ratio statistic

In Chapter 4, we explained how GstLAL assigns to each trigger a likelihood ratio defined by

$$\mathcal{L} = \frac{P(\vec{D}_H, \vec{O}, \vec{\rho}, \vec{\xi}^2, [\vec{\Delta}t, \vec{\Delta}\phi] | \vec{\theta}, \text{signal})}{P(\vec{D}_H, \vec{O}, \vec{\rho}, \vec{\xi}^2, [\vec{\Delta}t, \vec{\Delta}\phi] | \vec{\theta}, \text{noise})} \cdot \frac{P(\vec{\theta} | \text{signal})}{P(\vec{\theta} | \text{noise})}, \quad (5.1)$$

where (1) \vec{D}_H are the horizon distances of the detectors (a measure of the detector sensitivity), (2) \vec{O} is the set of detectors that registered the coincident trigger, (3) $\vec{\rho}$ are the recorded detector SNRs for the trigger at each participating detector, (4) $\vec{\xi}^2$ are the auto-correlation-based consistency test values evaluated for each of the participating detectors, (5) $[\vec{\Delta}t, \vec{\Delta}\phi]$ are the sets of differences in end times Δt and coalescence phases $\Delta\phi$ between pairs of detectors within the set of participating detectors and (6) $\vec{\theta}$ is the template parameters associated with the trigger.

The likelihood ratio can be factorized into products of smaller terms [85, 250]. This subsection will focus on the term $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$. This term gives the probability of obtaining a noise trigger with SNRs $\vec{\rho}$ and $\vec{\xi}^2$. To compute this probability, GstLAL populates 2D histograms in the ρ - ξ^2 space for each participating detector with non-coincident triggers, assuming they represent triggers originating from noise. The histograms are normalized and smoothed with a Gaussian smoothing kernel to produce an approximated $P(\vec{\rho}, \vec{\xi}^2 | \vec{\theta}, \text{noise})$ for assigning likelihood ratios for triggers. As noted in the previous chapter and [250], these histograms are constructed cumulatively throughout the data being analyzed, and no time dependence is tracked.

Consider the following scenario: The data quality is generally “bad” (i.e., the data are “noisy” and potentially contain numerous glitches) over a certain period of analysis time. We perform a search across the data that includes the aforementioned period of data with the GstLAL search pipeline. Glitches and noise in the noisy period of data get registered as noise triggers. These triggers, in particular those corresponding to glitches, may have SNRs ρ 's and ξ^2 very different from those of noise triggers from stationary and Gaussian noise. These triggers will enter the ρ - ξ^2 histograms and can potentially “pollute”¹ the expected probability distribution of ρ and ξ^2 for noise triggers due to their extreme values. At some later time, when the data becomes less “noisy”, a near-threshold, low-amplitude gravitational-wave signal may still get registered as a trigger. However, since the ρ - ξ^2 histogram for background triggers

¹ Generally speaking, the probability distribution is “polluted” if it is broadened by extreme data.

has been polluted by earlier noisy data, triggers with weaker SNRs may potentially be assigned a smaller likelihood ratio and significance, causing the pipeline to miss the signal potentially. In Section 5.3, we propose a method to potentially resolve this problem by modifying the likelihood ratio statistic to consider data quality around the time of the trigger when assigning ranking statistics.

The iDQ data quality pipeline: A brief introduction

iDQ is a machine-learning-based statistical inference pipeline that generates probabilistic information about the data quality in quasi-real time [149, 176]. The iDQ pipeline has been operating since LVK's first observation run. The main goal of the iDQ pipeline is to automatically identify potential non-Gaussian noise (glitches) in the data and labelled times in the data up to subsecond intervals as likely to contain a glitch in low latency. Details about the formalism, setup of the iDQ pipeline (including the training and training data generation), and effectiveness evaluation of the iDQ pipeline are all out of the scope of this thesis. Related information, however, can be found in [149]. One of the data products output from the iDQ pipeline is a log-likelihood estimate $\ln \mathcal{L}_{\text{iDQ}}$ of the presence of non-Gaussian noise in the data as a function of time (i.e. a timeseries. See Figure 5.2.1, for example). For practical reasons, as suggested in [176], $\ln \mathcal{L}_{\text{iDQ}}$ is renormalized to control the impact on search pipelines better when being incorporated into ranking statistic calculations. For a given duration of data being analyzed, the raw $\ln \mathcal{L}_{\text{iDQ}}$ timeseries sampled at 128 Hz is aggregated and maximized over a ± 1 second time window. The resulting values are then used to evaluate the renormalized iDQ log-likelihood defined by

$$\ln \hat{\mathcal{L}}_{\text{iDQ}} = \begin{cases} 0, & P < P_{\min}, \\ \ln \frac{P}{100-P}, & P_{\min} \leq P < P_{\max}, \\ 15, & P \geq P_{\max} \end{cases}, \quad (5.2)$$

where $P_{\min} = 50$, and

$$P_{\max} = \frac{100}{1 + \exp \left[-\ln \hat{\mathcal{L}}_{\text{iDQ}}^{\text{upper}} \right]} = \frac{100}{1 + e^{-15}}, \quad (5.3)$$

with $\ln \hat{\mathcal{L}}_{\text{iDQ}}^{\text{upper}}$ being the predefined upper limit of the renormalized iDQ log-likelihood, which is set to be 15. P is the percentile of $\ln \mathcal{L}_{\text{iDQ}}$ values across the whole timeseries. P_{\min} and P_{\max} are chosen accordingly to limit the renormalized $\ln \mathcal{L}_{\text{iDQ}}$ to a range between 0 and 15. Values below the 50th percentile are assigned a renormalized iDQ log-likelihood of 0. The higher $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ is, the more




Figure 5.2.1: An example timeseries of log-likelihood estimate $\ln \mathcal{L}_{\text{iDQ}}$ of the presence of non-Gaussian noise in the data output from the iDQ pipeline. The timeseries covers roughly 5000s of data (i.e. ≈ 1.3 hours). Note that $\ln \mathcal{L}_{\text{iDQ}}$ has a value of ≈ 9 most of time, but occasionally it can increase to 70 or more (e.g. see the peak at GPS times ≈ 2600 s). These peaks represent times the iDQ pipeline believes glitches are present in the data.

likely there exists non-Gaussian noise in the data. Interested readers are referred to [149, 176] for details.

5.3 Proposed method

Recall that the problem with the current implementation of the $P(\rho, \xi^2 | \text{noise})$ term in GstLAL’s log-likelihood ratio calculation is that non-Gaussianity in data can potentially “pollute” the cumulative ρ - ξ^2 histograms for background noise which are used to approximate $P(\rho, \xi^2 | \text{noise})$. The approximated, “polluted” $P(\rho, \xi^2 | \text{noise})$ function can lead to unnecessary reductions in the assigned log-likelihood ratio for triggers that are found in times where data quality is relatively good. This may lead to signal loss and a decrease in the search sensitivity of the pipeline.

Figure 5.3.1 shows a schematic overview of our proposed method to resolve the problem potentially. We propose to promote the current 2D probability density




Figure 5.3.1: The proposed scheme on modifying the original $P(\rho, \xi^2|\text{noise})$ term in GstLAL’s likelihood ratio calculations. We propose promoting the 2D ρ - ξ^2 histogram for noise triggers to 3D histograms such that the renormalized iDQ log-likelihood $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ as an additional parameter. In the figure, the proposed 3D histogram is represented by a series of slices of 2D ρ - ξ^2 histograms along the iDQ dimension. GstLAL will then populate the extended 3D histograms for each detector with noise triggers, according to the data quality around the time of each trigger informed by $\ln \hat{\mathcal{L}}_{\text{iDQ}}$. Then, as usual, the histograms will be normalized and smoothed with a Gaussian smoothing kernel. The end product will be a three-dimensional probability density function $P(\rho, \xi^2, \ln \hat{\mathcal{L}}_{\text{iDQ}}|\text{noise})$, which can be used to replace the original $P(\rho, \xi^2|\text{noise})$ term in the likelihood ratio calculation in GstLAL.

function $P(\rho, \xi^2|\text{noise})$ to a 3D probability density function $P(\rho, \xi^2, \ln \hat{\mathcal{L}}_{\text{iDQ}}|\text{noise})$ that includes the renormalized iDQ log-likelihood $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ as an additional parameter. The middle part of Figure 5.3.1 shows two “slices” of the proposed 3D probability density function, i.e. $P(\rho, \xi^2, \ln \hat{\mathcal{L}}_{\text{iDQ}}|\text{noise})$ at fixed iDQ log-likelihood values. The top and bottom histograms illustrate how the projected 2D probability density function will look like for low $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ (corresponding to times where data quality is good) and high $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ (corresponding to times where the data quality is bad) respectively. One will expect that the projection of the 3D probability density function $P(\rho, \xi^2, \ln \hat{\mathcal{L}}_{\text{iDQ}}|\text{noise})$ to span a larger region (i.e. the distribution is broader) in the ρ - ξ^2 space because it is populated and “polluted” with triggers coming from non-Gaussian features in the data (e.g. glitches).

In the modified scheme, GstLAL will set up 3D histograms in the SNR ρ - ξ^2 - $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ parameter space per detector, populated by noise triggers collected during the filtering process. These histograms, as usual, will then be normalized and smoothed with a Gaussian smoothing kernel to approximate the three-dimensional

probability density function $P(\rho, \xi^2, \ln \hat{\mathcal{L}}_{\text{iDQ}}|\text{noise})$, which is then used to assign likelihood ratios to signal triggers.

5.4 Expected results

Because of various reasons, implementing and testing the aforementioned method remains an ongoing work. I will briefly discuss the expected results. Figure 5.4.1

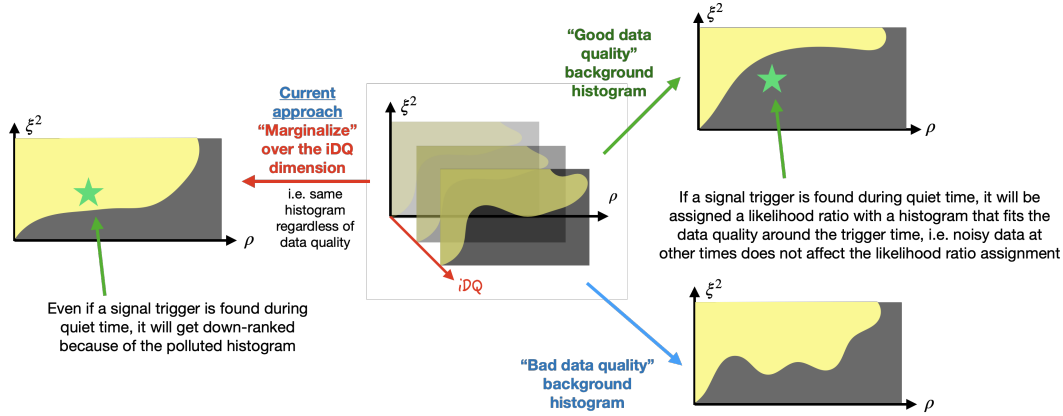


Figure 5.4.1: An illustration showing how the proposed method can improve the search sensitivity of GstLAL. Suppose we have a signal trigger found in a detector when the data quality is good. A green star in the figure represents the signal trigger. In the current GstLAL framework, since data quality information is not used, the 3D ρ - ξ^2 - $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ histograms (middle) for noise triggers are essentially “marginalized” over the data quality (iDQ) dimension, resulting in the usual 2D ρ - ξ^2 histogram on the left. As discussed before, non-Gaussianity (e.g. glitches) in the data can pollute the histograms and effectively broaden the probability distribution $P(\rho, \xi^2|\text{noise})$. As shown in the figure, the trigger found during times with good data quality lies in the region of the marginalized probability distribution $P(\rho, \xi^2|\text{noise})$ with high probability density, that is polluted by noise triggers associated with non-Gaussian noise in the data. Essentially, this means that $P(\rho, \xi^2|\text{noise})$ for this trigger will be unnecessarily high, resulting in a lower log-likelihood ratio and significance being assigned to it. In the proposed method, on the other hand, the trigger found in times with good data quality effectively has its log-likelihood ratios assigned using a histogram that suits the data quality around the time of the trigger (e.g. right upper histogram). As we can see, the trigger does not live within the probability distribution $P(\rho, \xi^2|\text{noise})$ for noise triggers coming from times with good data quality, and it will be assigned a lower $P(\rho, \xi^2|\text{noise})$ and hence higher log-likelihood ratio and significance. This helps to prevent unnecessary down-ranking of triggers, which may lead to potential loss of gravitational-wave signals.

illustrates how the proposed method can potentially improve the search sensitivity of GstLAL. Suppose we have a signal trigger found when the data quality is good. In the figure, the example trigger is represented by a green star. Under GstLAL’s

current framework, since data quality information is not considered, the 3D ρ - ξ^2 - $\ln \hat{\mathcal{L}}_{\text{iDQ}}$ histograms (in the middle of the figure) for noise triggers are essentially “marginalized” over the data quality (iDQ) dimension, resulting in the usual 2D ρ - ξ^2 histogram shown on the left. As discussed in Section 5.2, extreme noise triggers originating from non-Gaussianity (e.g. glitches) in the data can “pollute” the histograms and effectively broaden the probability distribution $P(\rho, \xi^2 | \text{noise})$. As shown in the left histogram in the figure, the trigger found during times with good data quality lies in the region of the marginalized probability distribution $P(\rho, \xi^2 | \text{noise})$ with high probability density that is “polluted” by extreme noise triggers associated with non-Gaussian noise in the data. Essentially, this means that $P(\rho, \xi^2 | \text{noise})$ for this trigger will be unnecessarily high, resulting in a lower log-likelihood ratio and significance being assigned to it. The trigger with a low assigned significance can potentially be buried in the noise background and not being identified as a gravitational wave signal, leading to a reduction in the search sensitivity of the pipeline.

In the proposed method, the trigger found in times with good data quality effectively has their log-likelihood ratios assigned using a histogram that suits the data quality around the time of the trigger (e.g. right upper histogram). As we can see, the trigger does not live within the probability distribution $P(\rho, \xi^2 | \text{noise})$ for noise triggers coming from times with good data quality, and it will be assigned a lower $P(\rho, \xi^2 | \text{noise})$ and hence higher log-likelihood ratio and significance. This helps to prevent unnecessary down-ranking of triggers that lead to potential loss of gravitational-wave signals.

Part II

Gravitational Lensing of Gravitational Waves

Chapter 6

INTRODUCTION TO GRAVITATIONAL LENSING AND LENSING OF GRAVITATIONAL WAVES

6.1 Overview

Part II of my thesis will focus on another phenomenon predicted in Einstein's theory of general relativity — Gravitational Lensing. This chapter is arranged as follows: In Section 6.2, I will first briefly introduce gravitational lensing. In Section 6.3, I will discuss gravitational lensing of gravitational waves. I will then introduce the different regimes of lensing of gravitational waves, i.e. microlensing, millilensing, and strong lensing. I will outline the current effort to look for lensing of gravitational waves. I will briefly introduce the different pipelines and technologies used by the LVK collaboration to search for the first lensed gravitational waves. Finally, in 6.4, I will describe briefly the content of the remaining chapters of this part of my thesis.

6.2 A very brief overview of gravitational lensing

Optical lensing

When we look at objects through an optical lens like a convex lens or concave lens, their images will appear to have different locations than the object's original position. They will also appear to have sizes (magnified/diminished) compared to the object. This can be understood as light is refracted when it passes through the optical lens. This is known as **optical lensing**. Figure 6.2.1 shows a typical example of how light rays are refracted under the influence of a convex lens.

The concept and history of gravitational lensing

Einstein's theory of general relativity describes the universe as a fabric of spacetime. Masses can produce curvature on this fabric of spacetime. When electromagnetic (EM) waves from a source pass by intervening massive objects like galaxies and galaxy clusters, their travelling paths will be deflected due to the curvature of spacetime. This is known as **gravitational lensing**. One interesting point is that the concept of light bending could already be seen in Newtonian mechanics before Einstein developed his theory of general relativity. I invite interested readers to the Appendix Section 6.A.

An important concept in gravitational lensing is the **deflection angle**. The deflection




Figure 6.2.1: A cartoon showing a typical case of optical lensing with a convex lens. Light from the object on the left is being focused onto the observer on the right by the lens. The observer does not know that the light has been refracted and will think that it originates from an image (in pink) at a further position (different from the object's position), with a larger size (magnified).

angle measures how much deflection a gravitational lens induces on the travelling paths of the waves emitted from a source. Since light bending is already predicted in Newtonian mechanics, it should not be surprising that the deflection angle was calculated way before Einstein formulated his theory of general relativity. Henry Cavendish was the first to calculate the gravitational lensing deflection angle back in 1784, followed by Johann Von Soldner in 1801 [146]. In 1911, Einstein also attempted to evaluate the deflection angle using Newtonian mechanics, and the deflection angle near the surface of the Sun was determined to be roughly $\delta\theta = 0.84''$ [133]. In 1913, Einstein wrote a letter to astronomer George Hale and asked if it was possible to observe and measure a deflection angle of $\delta\theta = 0.84''$. George Hale explained to Einstein that unless it is in the case of a solar eclipse, the background starlight will be dominated by sunlight. In 1914, Einstein asked Erwin Freundlich to lead a German team to Russia to observe a total solar eclipse to verify the deflection angle calculation [141]. However, the team was captured during the First World War (WWI) and was released only after the solar eclipse.

In 1915, Einstein recalculated the deflection angle with correction based on general relativity [174] and obtained $\delta\theta = 1.75''$. The detailed calculations for the relativistic deflection angle are found in the Appendix Section 6.C. Sir Eddington set out to observe a total solar Eclipse on May 29, 1919, partially to avoid military service that would conflict with his Quaker faith [141, 173]. Eddington observed a deflection angle of $\delta\theta = 1.9''$, proving Einstein was right up to a 30% error. The deflection

of light is the first confirmed prediction of general relativity, even before the proper explanation of the perihelion shift of Mercury [111].

Unexpectedly, Einstein did not believe gravitational lensing would be a significant contribution [147]. This is because he only considered stellar-mass objects as possible gravitational lenses, for which the images will not be observationally resolvable. In 1937, Fritz Zwicky considered the possibility of nebulae being gravitational lenses [378]. He also considered the possibility of a galaxy being observed as multiple distinct lensed images if its light passed sufficiently close to a foreground cluster of galaxies[111]. The foreground cluster of galaxies has to contain a lot of unseen dark matter to prevent the galaxies from flying apart. In 1970, Vera Rubin and Kent Ford demonstrated that a single galaxy could also act as a gravitational lens [307].

Gravitational lensing was long forgotten until the 1960s when Norwegian Sjur Refsdal started working on it again. In 1964, Sjur proposed that measurements of the Hubble constant can be made independent of other methods by measuring the time delay between lensed images [208]. People in the (astro)physics community were initially skeptical, but later decided it was a plausible idea. Since then, gravitational lensing has become a highly useful tool in astrophysics. For instance, gravitational lensing can be used to search for the presence of dark matter [107, 245] and exoplanets [70], and to detect and discover faint massive objects and structures [108].

Key concepts in gravitational lensing

I will now summarize the key concepts in gravitational lensing. Refer to Figure 6.2.2 for the discussion.

In gravitational lensing, we focus on the “lens equation”, given by

$$\vec{\beta} + \frac{D_{LS}}{D_S} \hat{\alpha} = \vec{\theta}. \quad (6.1)$$

The lens equation relates the source position $\vec{\beta}$, the image position $\vec{\theta}$, the source distance D_S , the lens-source distance D_{LS} and the deflection angle $\vec{\alpha}$ together. Appendix section 6.D provides a detailed geometric derivation for the lens equation. One can also derive the lens equation by considering the total time delay of waves emitted from a source arriving at the observer due to gravitational lensing. There




Figure 6.2.2: Parameters for deriving the lens equation. Here, the solid black line joining the observer and the center of the lens defined the z -axis. All angles described in this figure are two-dimensional azimuthal angles (e.g. one can decompose the angles into x and y components). $\vec{\theta}$ is the 2D angle of the image position. $\vec{\xi}$ is the 2D impact parameter. $\vec{\eta}$ is the 2D source position. $\vec{\beta}$ is the 2D angle of the source position. All other parameters have the same meanings as in Figure 6.D.2.

are two components in the time delay function, namely, the geometric time delay

$$\tau_{\text{geo}} = (1 + z_L)\Delta\lambda = (1 + z_L)\frac{D_S D_L}{2D_{LS}}|\vec{\theta} - \vec{\beta}|^2, \quad (6.2)$$

where z_L is the lens redshift; and gravitational time delay (also known as the Shapiro time delay)

$$\tau_{\text{grav}} = -2(1 + z_L) \int \Phi dl \quad (6.3)$$

where Φ is the Newtonian potential of the lens. Detailed derivation of the two components can be found in the Appendix section 6.E.

The total time delay caused by gravitational lensing is given by

$$\tau = \tau_{\text{geo}} + \tau_{\text{grav}} \quad (6.4)$$

$$= (1 + z_L)\frac{D_S D_L}{2D_{LS}}|\vec{\theta} - \vec{\beta}|^2 - 2(1 + z_L) \int \Phi dl \quad (6.5)$$

$$= (1 + z_L)\frac{D_S D_L}{D_{LS}} \left[\frac{1}{2}|\vec{\theta} - \vec{\beta}|^2 - \Phi_{\text{eff}} \right], \quad (6.6)$$

where

$$\Phi_{\text{eff}} = \frac{D_{LS}}{D_L D_S} \int 2\Phi dl \quad (6.7)$$

is the “effective 2D potential” of the lens. The term inside the bracket is known as the Fermat potential. According to the Fermat principle, light travels along paths with stationary traveling times. By differentiating the Fermat potential with respect to $\vec{\theta}$ and setting the results to be zero, i.e.

$$\nabla_{\vec{\theta}} \left[\frac{1}{2} |\vec{\theta} - \vec{\beta}|^2 - \Phi_{\text{eff}} \right] = 0, \quad (6.8)$$

we obtain

$$\vec{\theta} - \vec{\beta} - \nabla \Phi_{\text{eff}} = 0, \quad (6.9)$$

which is the same lens equation we have obtained previously, if we identify $\nabla \Phi_{\text{eff}}$ as the deflection angle, i.e.

$$\vec{\alpha} = \nabla \Phi_{\text{eff}}. \quad (6.10)$$

In practice, light will travel along paths where $\nabla \tau = 0$. To determine the type of solutions, we take the second derivative of the time delay function, which gives the lensing Jacobian:

$$\nabla \nabla \tau = \begin{bmatrix} \frac{d^2 \tau}{d\theta_x^2} & \frac{d^2 \tau}{d\theta_x d\theta_y} \\ \frac{d^2 \tau}{d\theta_x d\theta_y} & \frac{d^2 \tau}{d\theta_y^2} \end{bmatrix}. \quad (6.11)$$

The solution is a local minimum if the second derivative yields two positive eigenvalues. If it yields one positive and one negative eigenvalue, the solution is a saddle point solution. If it yields two negative eigenvalues, the solution is a local maximum.

In index notation form, the lens equation can be written as

$$\tau_{,i} = \theta_i - \beta_i - \psi_{,i}. \quad (6.12)$$

Taking the second derivative, we get

$$\tau_{,ij} = \underbrace{\theta_{i,j}}_{=\delta_{ij}} - \underbrace{\beta_{i,j}}_{=0} - \psi_{,ij} = \delta_{ij} - \psi_{,ij}, \quad (6.13)$$

where $\theta_{i,j} = \partial_j \theta_i$. Hence, we can write the lensing Jacobian as

$$A = \tau_{,ij} = \begin{bmatrix} 1 - \psi_{,xx} & -\psi_{,xy} \\ -\psi_{,xy} & 1 - \psi_{,yy} \end{bmatrix}. \quad (6.14)$$

The convergence is related to the potential by $2\kappa = \psi_{,ii}$, so we can write

$$\kappa = \frac{1}{2} [\psi_{,xx} + \psi_{,yy}]. \quad (6.15)$$

Hence, the lensing Jacobian can be rewritten as

$$A = \tau_{,ij} = \begin{bmatrix} 1 - \kappa - \frac{1}{2} [\psi_{,xx} - \psi_{,yy}] & -\psi_{,xy} \\ -\psi_{,xy} & 1 - \kappa + \frac{1}{2} [\psi_{,xx} - \psi_{,yy}] \end{bmatrix} \quad (6.16)$$

$$= \begin{bmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_2 \end{bmatrix}, \quad (6.17)$$

where

$$\begin{cases} \gamma_1 = \frac{1}{2} [\psi_{,xx} - \psi_{,yy}] \\ \gamma_2 = \psi_{,xy} \end{cases}. \quad (6.18)$$

We can write the Jacobian into a sum of two matrices:

$$A = \tau_{,ij} = (1 - \kappa) \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - \begin{bmatrix} \gamma_1 & \gamma_2 \\ \gamma_2 & -\gamma_1 \end{bmatrix}. \quad (6.19)$$

The first matrix corresponds to an isotropic scaling of the image, whereas the second matrix corresponds to the shearing of the image.

To find the magnification of the image, we can evaluate the magnification matrix

$$M = \frac{d\vec{\theta}}{d\vec{\beta}}. \quad (6.20)$$

The magnitude of the magnification is simply given by the determinant of the magnification matrix, i.e.

$$\mu = |M|. \quad (6.21)$$

From the lens equation, we have

$$\beta_{i,j} = \theta_{i,j} - \psi_{,ij} = A_{ij}. \quad (6.22)$$

This means that the lensing Jacobian is simply the inverse of the magnification matrix, i.e. $A = M^{-1}$.

There exist positions where the magnification becomes infinite. The corresponding curves $\vec{\theta}(\mu \rightarrow \infty)$ on the lens plane are known as critical curves. We can also map these curves back to the image plane using the lens equation. These curves




Figure 6.2.3: Critical curves and caustics of a point mass lens model.

are known as the caustics. Figure 6.2.3 shows the caustics and critical curves of a point mass lens model. The caustics also provide a very intuitive way to determine the number of images formed. A new pair of images will be formed whenever the source crosses a caustic.

For a more detailed discussion on gravitational lensing in general, readers are referred to standard reviews and textbooks, including but not limited to [111, 133, 141, 315]. I will end this section here by stating that gravitational lensing has been extensively studied with electromagnetic waves [52, 70, 107, 108, 117, 142, 245, 251, 271, 299, 348]. Starting from the next section, I will focus on discussing gravitational lensing of gravitational waves.

6.3 Gravitational lensing of gravitational waves

In this section, I will focus on the lensing of gravitational waves. I will first explain some of the motivations for studying lensed gravitational waves. Then, I will briefly discuss the different regimes of gravitational lensing of gravitational waves, from microlensing to millilensing, and end with strong lensing. For the rest of my thesis, I will focus on strong lensing of gravitational waves. Next, I will discuss the current effort from the LVK collaboration to search for lensing signatures in gravitational-wave data, outlining the key pipelines and technologies being used in preparation for discussions in the later chapters.

Motivations for studying lensed gravitational waves

The effect of gravitational lensing has been exclusively studied with electromagnetic waves. Gravitational waves should be gravitationally lensed in the same way, according to the equivalence principle. As discussed in [23, 32, 186, 203], the

search for gravitationally-lensed gravitational waves is highly motivated.

One of the avenues is to test general relativity [51, 110, 154, 155, 178, 179]. Suppose we identify a pair of strongly lensed gravitational waves originating from the same source with a network of three linearly independent detectors (e.g. the LIGO-Hanford, LIGO-Livingston, and Virgo detectors). We are essentially observing the same signal twice, albeit at different times. An alternative view is that we are observing the gravitational-wave signal with a six-detector network. This provides an excellent opportunity to perform precise polarization tests on gravitational waves and test alternative theories of gravity [124, 178, 241] as the network can measure different combinations of the same polarization. One may also test the speed of gravitational waves over cosmological distances with strongly lensed gravitational waves (See [110]).

Another possible avenue is to perform precision cosmology [88, 187, 225, 227, 317] and statistical cosmology [372]. Strongly lensed gravitational waves have accurately measured time delays. With the help of multi-messenger astronomy, we can observe images and redshifts in the electromagnetic regime. Together, the two pieces of information allow us to achieve precise cosmography and a better-measured Hubble constant [197, 227]. Detecting strongly lensed gravitational waves also allows us to probe the population of compact binaries and galaxies, especially in the higher redshift universe [372]. Understanding the early universe through gravitational waves is not possible unless one takes into account gravitational lensing of gravitational waves, and the stochastic gravitational-wave background.

Studying and detecting gravitationally-lensed gravitational waves also allows one to improve source localization [187, 374] and lens model characterization [138, 220, 230, 283]. Apart from traditional candidates for gravitational lenses, including galaxies and galaxy clusters, one can also make use of gravitationally-lensed gravitational waves to probe dark matter in the universe, an alternative viable candidate for gravitational lenses [89, 158, 231, 376], and possibly halo properties [339].

Basics of lensing of gravitational waves

The derivations and discussion in this subsection closely follow [154, 259]. In Section 6.2, we show that the total time delay between the lensed and unlensed travelling paths in the thin lens approximation is given by the sum of the geometric time delay τ_{geo} and the gravitational time delay (also known as the Shapiro time

delay) τ_{grav} :

$$\tau(\vec{\theta}, \vec{\beta}) = \tau_{\text{geo}} + \tau_{\text{grav}} \quad (6.23)$$

$$= (1 + z_L) \frac{D_S D_L}{2 D_{LS}} |\vec{\theta} - \vec{\beta}|^2 - 2(1 + z_L) \int \Phi dl \quad (6.24)$$

$$= (1 + z_L) \frac{D_S D_L}{D_{LS}} \left[\frac{1}{2} |\vec{\theta} - \vec{\beta}|^2 - \Phi_{\text{eff}} \right], \quad (6.25)$$

where

$$\Phi_{\text{eff}} = \frac{D_{LS}}{D_L D_S} \int 2\Phi dl \quad (6.26)$$

is the effective gravitational potential.

Following [154], the observed lensed gravitational wave is given by

$$\tilde{h}_L = \frac{1}{2\pi} \int F(\omega, \vec{\beta}) \cdot \tilde{h}_U d\omega, \quad (6.27)$$

where ω is the angular frequency,

$$\tilde{h}_U = \int h(t) e^{i\omega t} dt \quad (6.28)$$

is the frequency domain unlensed gravitational wave, and the amplification factor

$$F(\omega, \vec{\beta}) = \frac{D_L D_S}{D_{LS}} \frac{\omega}{2\pi i} \int e^{i\omega\tau(\vec{\theta}, \vec{\beta})} d^2\theta \quad (6.29)$$

considers all the possible traveling paths of the gravitational waves from the source to the observer. The amplification factor is also known as the **Fresnel-Kirchhoff diffraction integral**, a reduction of the infinite-dimensional Feynman path integral [154, 269]. Note that the amplification factor in the most general case is **frequency-dependent**.

Depending on the mass properties of the lens [134, 270, 284, 337, 340, 361], the effect of gravitational lensing can be vastly different, and different approximations can also be applied. In the following sections, we will investigate the various regimes of gravitational lensing: microlensing, millilensing, and strong lensing.

Microlensing of gravitational waves

When we consider gravitational lenses with sizes that are comparable to gravitational-wave wavelengths, they can produce frequency-dependent modulations, also known as wave optics effects, on the lensed waveforms via the amplification factor array

(See equation 6.29). This is formally known as **microlensing** [90, 93, 95, 101, 105, 116, 121, 134, 138, 220, 248, 256, 270, 337]. The detection of microlensed gravitational waves allow us to study and understand the nature of the gravitational lens [90, 101, 105, 121, 138, 139, 205, 220, 248, 257, 286, 337].

Gravitational lenses with masses $\leq O(10^5)M_\odot$ fall within the microlensing regime. These include stellar-mass objects and intermediate-mass black holes. Note that, however, it is not likely that a gravitational wave detected on the Earth only passes by a single stand-alone microlens, but rather, these microlenses are often embedded within larger structures such as a galaxy or galaxy clusters. Therefore, more realistic models where one or more microlenses are included in a larger macrolens have been considered and developed [101, 139, 257, 333, 373], albeit resulting in more complicated waveforms. Nevertheless, due to the expensive computational costs to work with these realistic models, the LVK collaboration has only been searching for microlensing effects in gravitational waves assuming an isolated lens [23, 32, 186, 203, 316] (See also the subsection “Current Effort from LVK Collaboration”) with different lens model, including the Point Mass Lens model (PML), the Singular Isothermal Sphere model (SIS), the Singular Isothermal Ellipsoid model (SIE), and the Navarro-Frenk-White model (NFW).

Strong Lensing of gravitational waves

When the size of the gravitational lenses is (much) larger than the gravitational-wave wavelength, the time delay between stationary paths (among all the possible travelling paths) in the amplification factor (path) integral becomes (much) larger than the inverse frequency of the wave, and hence we can approximate the integral by assuming that it is dominated by the stationary phase points in the Fermat’s time-delay surface under the stationary phase approximation [154, 259]. These stationary phase points are essentially unique separate lensed gravitational-wave signals (different from the inseparable signals from microlensing). The Fermat’s time-delay surface allows for three stationary point solutions: a minimum, a saddle-point and a maximum. The lensed images formed at these points are referred to as Type I, Type II and Type III images accordingly.

Following [154], we can Taylor-expand the total time delay τ around the stationary phase points $\vec{\theta}_k$ under the stationary phase approximation as

$$\tau(\vec{\theta}, \vec{\beta}) \approx \tau(\vec{\theta}_k) + \frac{1}{2} \sum_{(a,b)=1}^2 |\vec{\theta}_k - \vec{\beta}|_a |\vec{\theta}_k - \vec{\beta}|_b \partial_a \partial_b \tau(\vec{\theta}_k), \quad (6.30)$$

where k labels the k^{th} image, and a, b represents the two components of the two-dimensional angles. Note that the first-order term in the Taylor expansion vanishes because the time-delay function is expanded around the **stationary phase points**. From the total time delay function $\tau(\vec{\theta}, \vec{\beta})$, we can obtain the stationary points by solving

$$\partial_a \tau(\vec{\theta}) = 0. \quad (6.31)$$

The nature of each of these stationary points, as we have seen before, can be determined by evaluating the second-order derivative of the total time delay function at these positions. The result would be a 2×2 matrix known as the lensing Jacobian (or the Hessian matrix [154])

$$T_{ab} = \partial_a \partial_b \tau(\vec{\theta}) = \begin{bmatrix} \frac{d^2 \tau}{d\theta_x^2} & \frac{d^2 \tau}{d\theta_x d\theta_y} \\ \frac{d^2 \tau}{d\theta_x d\theta_y} & \frac{d^2 \tau}{d\theta_y^2} \end{bmatrix}. \quad (6.32)$$

The stationary point is a local minimum if the second derivative yields two positive eigenvalues. If it yields one positive and one negative eigenvalue, the solution is a saddle point solution. If it yields two negative eigenvalues, the solution is a local maximum. We can diagonalize the lensing Jacobian in terms of the eigenvalues $\lambda_{1k}, \lambda_{2k}$ and corresponding eigenvectors $\tilde{\theta}_1, \tilde{\theta}_2$. This allows us to write the second term in the Taylor expansion for the total time delay function as

$$\frac{1}{2} \sum_{(a,b)=1}^2 |\vec{\theta}_k - \vec{\beta}|_a |\vec{\theta}_k - \vec{\beta}|_b \partial_a \partial_b \tau(\vec{\theta}_k) = \frac{1}{2} \frac{D_S D_L}{D_{LS}} \left(\lambda_{1k} \tilde{\theta}_1^2 + \lambda_{2k} \tilde{\theta}_2^2 \right). \quad (6.33)$$

Substituting this into the amplification factor array, we have

$$F = \frac{D_L D_S}{D_{LS}} \frac{\omega}{2\pi i} \int e^{i\omega \tau(\vec{\theta}, \vec{\beta})} d^2 \theta \quad (6.34)$$

$$\approx \frac{D_L D_S}{D_{LS}} \frac{\omega}{2\pi i} \times \sum_k e^{i\omega \tau(\vec{\theta}_k)} \times \int \exp \left[\frac{i\omega}{2} \frac{D_S D_L}{D_{LS}} \left(\lambda_{1k} \tilde{\theta}_1^2 + \lambda_{2k} \tilde{\theta}_2^2 \right) \right] d^2 \tilde{\theta} \quad (6.35)$$

$$= \frac{D_L D_S}{D_{LS}} \frac{\omega}{2\pi i} \times \sum_k e^{i\omega \tau(\vec{\theta}_k)} \times \int \exp \left[\frac{i\omega \lambda_{1k}}{2} \frac{D_S D_L}{D_{LS}} \tilde{\theta}_1^2 \right] d\tilde{\theta}_1 \quad (6.36)$$

$$\times \int \exp \left[\frac{i\omega \lambda_{2k}}{2} \frac{D_S D_L}{D_{LS}} \tilde{\theta}_2^2 \right] d\tilde{\theta}_2. \quad (6.37)$$

Note that the last two terms are complex Gaussian integrals. We can evaluate the integrals with the identity

$$\int_{-\infty}^{\infty} \exp [\pm i x^2] dx = \exp \left[\pm i \frac{\pi}{4} \right] \cdot \sqrt{\pi}. \quad (6.38)$$

For the first integral, we make the following change of variable (in [154], it was phrased as a renormalization)

$$\tilde{\chi}_1 = \sqrt{\frac{\omega \lambda_{2k}}{2} \frac{D_S D_L}{D_{LS}}} \tilde{\theta}_1, \quad (6.39)$$

which transforms the integral as

$$\int \exp \left[\frac{i\omega \lambda_{1k}}{2} \frac{D_S D_L}{D_{LS}} \tilde{\theta}_1^2 \right] d\tilde{\theta}_1 = \sqrt{\frac{2}{\omega \lambda_{1k}} \frac{D_{LS}}{D_S D_L}} \int \exp \left[i \text{sign}(\omega \lambda_{1k}) \tilde{\chi}_1^2 \right] d\tilde{\chi}_1 \quad (6.40)$$

$$= \sqrt{\frac{2}{\omega \lambda_{1k}} \frac{D_{LS}}{D_S D_L}} \cdot \sqrt{\pi} \cdot \exp \left[i \text{sign}(\omega \lambda_{1k}) \frac{\pi}{4} \right]. \quad (6.41)$$

Note that the $\text{sign}(\dots)$ function denoting the sign of the argument (in this case, it is the product of the angular frequency ω and the eigenvalue λ_{1k}) is required as it will affect the evaluation of the Gaussian integral. The second integral can be transformed similarly into

$$\int \exp \left[\frac{i\omega \lambda_{2k}}{2} \frac{D_S D_L}{D_{LS}} \tilde{\theta}_2^2 \right] d\tilde{\theta}_2 = \sqrt{\frac{2}{\omega \lambda_{2k}} \frac{D_{LS}}{D_S D_L}} \cdot \sqrt{\pi} \cdot \exp \left[i \text{sign}(\omega \lambda_{2k}) \frac{\pi}{4} \right]. \quad (6.42)$$

Then, the amplification factor array can be written as

$$F \approx \frac{D_L D_S}{D_{LS}} \frac{\omega}{2\pi i} \times \sum_k e^{i\omega \tau(\vec{\theta}_k)} \times \frac{2\pi}{\omega} \frac{D_{LS}}{D_S D_L} \times \sqrt{\lambda_{1k} \lambda_{2k}} \quad (6.43)$$

$$\times \exp \left[i \left[\text{sign}(\omega \lambda_{1k}) + \text{sign}(\omega \lambda_{2k}) \right] \frac{\pi}{4} \right] \quad (6.44)$$

$$= \sum_k (-i) \cdot \exp \left[i\omega \tau(\vec{\theta}_k) \right] \cdot \sqrt{\frac{1}{|\lambda_{1k} \lambda_{2k}|}} \cdot \exp \left[i \text{sign}(\omega) \frac{q\pi}{2} \right], \quad (6.45)$$

where

$$q = \begin{cases} 1, & \text{for local minimum (two positive eigenvalues)} \\ 0, & \text{for saddle point (one positive and one negative eigenvalue)} \\ -1, & \text{for local minimum (two negative eigenvalues)} \end{cases} \quad (6.46)$$

(I use q to distinguish the usual n in most literature). We can write $-i = \exp \left[-i\frac{\pi}{2} \right]$, and hence the amplification factor array as

$$F = \sum_k \underbrace{\sqrt{\frac{1}{|\lambda_{1k}\lambda_{2k}|}}}_{\sqrt{|\mu(\vec{\theta}_k)|}} \cdot \exp \left[i\omega\tau(\vec{\theta}_k) - i\frac{\pi}{2} + i\text{sign}(\omega)\frac{q\pi}{2} \right] \quad (6.47)$$

$$= \sum_k \sqrt{|\mu(\vec{\theta}_k)|} \cdot \exp \left[i\omega\tau(\vec{\theta}_k) - i\text{sign}(\omega)\frac{n\pi}{2} \right], \quad (6.48)$$

where

$$\mu(\vec{\theta}_k) = \frac{1}{|\lambda_{1k}\lambda_{2k}|} \quad (6.49)$$

is the amplification factor for the k^{th} image (note that we have shown previously that the magnification is identical to the inverse of the product of the eigenvalues of the Lensing Jacobian), and

$$n = \begin{cases} 0, & \text{for local minimum (Type 1)} \\ 1, & \text{for saddle point (Type 2)} \\ 2, & \text{for local minimum (Type 3)} \end{cases}, \quad (6.50)$$

is commonly known as the **Morse phase**. Figure 6.3.1 illustrates the three types of strongly-lensed gravitational waves with different Morse phases.

Different lens models predict different distributions of relative time delays and relative magnifications between gravitationally-lensed gravitational waves from the same source. Gravitational lenses with moderate masses (e.g. galaxies) can produce lensed gravitational waves with relative time delays ranging from days to months, whereas more massive gravitational lenses (e.g. galaxy clusters) can produce lensed signals with relative time delays ranging from months to years [259, 281]. If appropriately used, the aforementioned prior information can help us to estimate the lensing rates for gravitational waves (See Chapter 8 and [266]) and can also help us to improve the search sensitivity for possible strongly-lensed gravitational waves (See Chapter 8).

It is also important to note that the magnification factor $\sqrt{|\mu_j|}$ can be smaller than 1 [281]. Hence, some strongly-lensed gravitational waves may be de-magnified with a reduced signal-to-noise (SNR) ratio below the detection threshold. These signals, formally known as “sub-threshold” signals, are of particular interest, as a significant

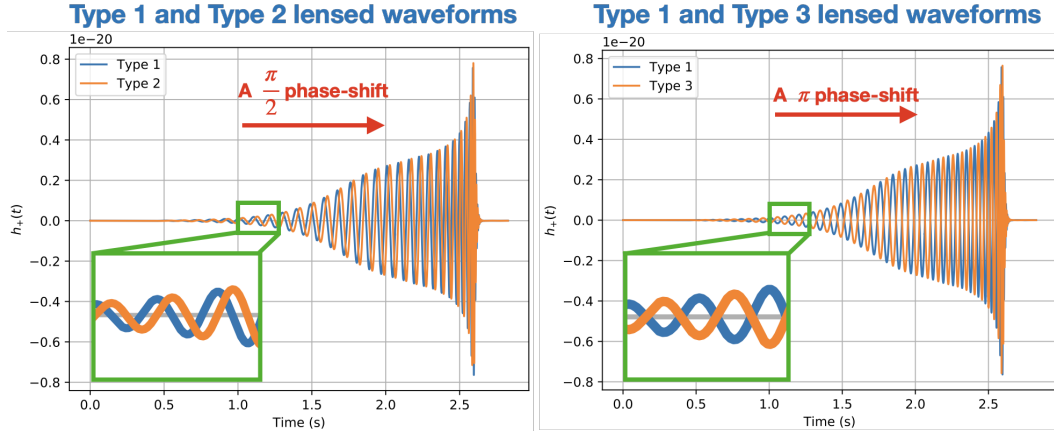


Figure 6.3.1: (Left) Type 1 (blue) and Type 2 (orange) gravitationally-lensed gravitational waveforms in the time domain. (Right) Type 1 (blue) and Type 3 (orange) gravitationally-lensed gravitational waveforms in the time domain. Type 1 signals have zero Morse phase shift, and hence, they are in phase with the not-lensed gravitational wave. Type 2 signals have a $\frac{\pi}{2}$ Morse phase shift compared to the not-lensed gravitational wave. Type 3 signals have a π Morse phase shift compared to the not-lensed gravitational wave.

proportion of strongly-lensed gravitational-wave events falls into this category [364]. Retrieving these possible sub-threshold lensed counterparts is a crucial step to boost the confidence for the first detection of strongly-lensed gravitational waves.

To summarize, in the strong lensing regime, we can use the stationary phase approximation to simplify the amplification factor integral so that it becomes frequency-independent (i.e. achromatic). This approximation is also known as the geometric optics limit. Mathematically, if we denote the not-lensed gravitational wave in the frequency domain as $\tilde{h}^{\text{NL}}(f; \vec{\theta}, \Delta t_j = 0)$, then the j^{th} strongly-lensed counterpart waveform \tilde{h}_j is

$$\tilde{h}_j(f; \vec{\theta}, \mu_j, \Delta t_j, \Delta \phi_j) = \sqrt{|\mu_j|} \tilde{h}^{\text{NL}}(f; \vec{\theta}, \Delta t_j) e^{(i \text{sign}(f) \Delta \phi_j)}. \quad (6.51)$$

where

$$\Delta \phi_j = -\frac{n_j \pi}{2}, \quad n_j = \begin{cases} 0, & \text{Type 1 lensed signals} \\ 1, & \text{Type 2 lensed signals} \\ 2, & \text{Type 3 lensed signals} \end{cases} \quad (6.52)$$

is the Morse phase factor, and μ_j is the magnification for the j^{th} lensed signal.

Millilensing of gravitational waves

In between strong lensing and microlensing, there is an intermediate regime where the sizes of the gravitational lenses are comparable to, although still larger than, the gravitational-wave lengths. In this regime, the geometric optics limit used for strong lensing is still valid, such that the lensed signals are still distinct signals. However, the relative time delays between the lensed signals are shortened so that they can start overlapping and interfere with each other, resulting in similar beating patterns as in microlensing. This is known as millilensing [230]. We expect millilensing to occur for gravitational lenses with masses ranging from $10^2 M_\odot$ to $10^6 M_\odot$. In this regime, the millilensed gravitational-wave signal is the sum of all the lensed gravitational-wave signals produced

$$\tilde{h}(f; \vec{\theta}, \{\mu_j, t_j, n_j\}_{j=1, \dots, K}) = \left[\sum_{j=1}^K \sqrt{|\mu_j|} e^{(i \text{sign}(f) \Delta \phi_j)} \right] \tilde{h}^{\text{NL}}(f; \vec{\theta}, \Delta t_j). \quad (6.53)$$

where j is the j^{th} lensed gravitational-wave signals, running from 1 to K . The interference between the lensed gravitational-wave signals results in a non-trivial millilensed gravitational wave. Since the total number of lensed signals is usually not known a priori, searching for such millilensed signals requires one to either fix the number of lensed signals in a search, or make it into a variable and try to infer it [203, 230].

Current LVK effort to search for lensing signatures in gravitational-wave data

Since the first successful detection of gravitational waves by the LIGO-Virgo-KAGRA (LVK) Collaboration [5, 6, 8, 11, 12, 17–19, 21, 24, 33], The LVK collaboration has conducted various analyses searching for lensing signatures in the event catalogs of LVK’s observing runs O1-O2 [186] and O3 [23, 32]. The details will be covered in later chapters. In this subsection, I will briefly introduce the framework of LVK’s lensing searches.

Figure 6.3.2 illustrates the general framework of the LVK collaboration’s lensing searches. We start with a set of confirmed gravitational-wave events documented in LVK’s Gravitational-Wave Transient Catalogues (GWTC). For each event in the set, we first perform microlensing analysis to look for possible microlensing signatures embedded in the gravitational-wave signal via the pipeline GRAVELAMPS [370]. In the follow-up work [203] for the full-O3 lensing analysis [32], we also search for possible millilensing signatures with the methodology described in [230]. LVK also searches for possible strongly-lensed gravitational waves within the available data.

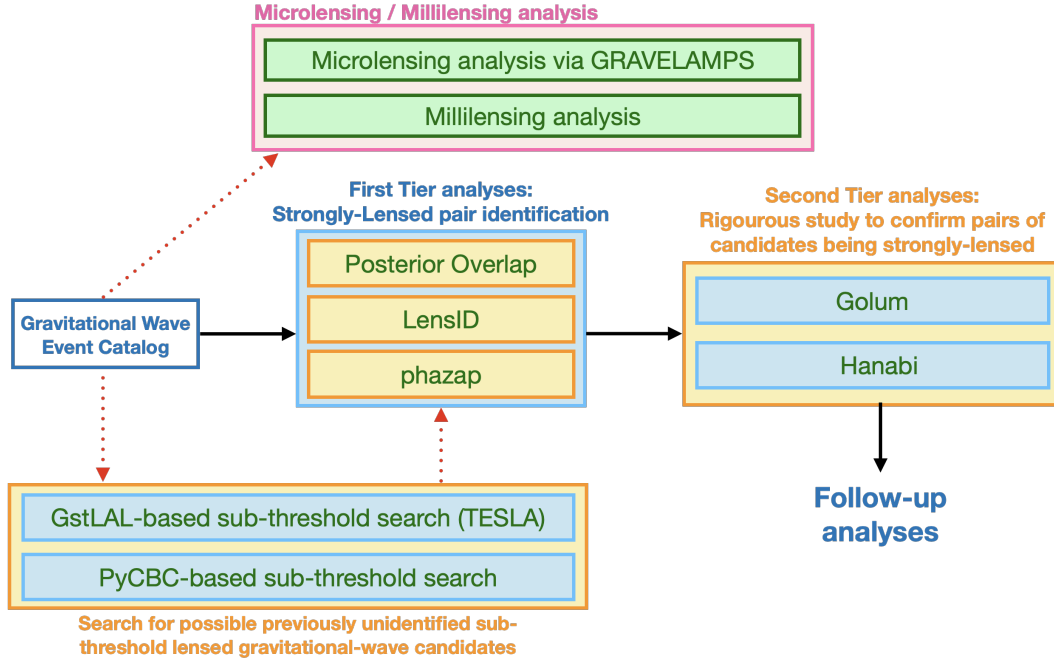


Figure 6.3.2: The general framework of the LIGO-Virgo-KAGRA (LVK) collaboration’s lensing searches. We start with a set of confirmed gravitational-wave events. We perform microlensing (via the pipeline GRAVELAMPS) and millilensing analysis for each event. We also search for possible sub-threshold lensed counterparts to each event via two independent search methods: (1) GstLAL-based search pipeline TESLA, and (2) PyCBC-based sub-threshold search. From the final pool of gravitational-wave candidates, we form pairs of candidates. For each pair, we utilize three independent preliminary tests/pipelines to check for signal consistency, namely (1) posterior overlap (checking for consistency in signal source parameters), (2) machine-learning-based classification test LensID (checking for consistency in sky location of source and time-frequency maps of signals), and (3) phazap (checking for consistency in the signal phase with the strong lensing hypothesis). Each of these tests assigns a ranking statistic to the pairs of candidates, and those that satisfy a pre-determined threshold are passed on for two independent, robust tests to determine how likely they are strongly-lensed counterparts to each other from the same source, namely (1) Golum and (2) Hanabi. Pairs that survive the tests with strong evidence of being strongly lensed will then be followed up further.

To account for possible de-magnified sub-threshold strongly-lensed gravitational waves that are not previously identified as a signal, we first employ two independent search pipelines: (1) GstLAL-based search pipeline TESLA [222, 223], and (2) PyCBC-based sub-threshold search [247] to search for possible sub-threshold lensed counterparts to known gravitational-wave events in the set under consideration. Then, we form possible pairs of lensed gravitational-wave candidates from the final pool of gravitational waves. For each pair, we utilize independent preliminary

tests/pipelines to check for signal consistency ¹:

- (1) Posterior overlap: Check for consistency of intrinsic parameters of the gravitational-wave signals by comparing the posteriors of gravitational waves obtained from standard parameter estimation under the unlensed hypothesis [188].
- (2) LensID: A machine-learning-based pipeline that checks for consistency in the sky location of the gravitational-wave sources and time-frequency maps of the signals [180].
- (3) phazap: Employed since the O4 analysis that checks for consistency of the signals' phase with respect to the strong lensing hypothesis (under strong lensing, only certain integral values of phase differences between lensed signals are allowed) [157].

Each of these pipelines assigns a ranking statistic to the pairs of possible lensed gravitational-wave candidates, and those that satisfy a pre-determined threshold are passed on to two independent, robust joint parameter estimation pipelines, Golum [200, 202] and Hanabi [235], to re-analyse the event pair under the strong-lensing hypothesis. At the end, the two pipelines evaluate a Bayes factor for each event pair, comparing the hypothesis that the two events are strongly-lensed counterparts to each other, and that the two events are independent, unlensed events. Pairs with sufficiently high evidence of being strongly lensed will then be followed up further. Details of these tests can be found in their respective references (cited above), as well as in later chapters.

6.4 Overview of the upcoming chapters

In Chapter 7, I will introduce the GstLAL-based search pipeline TESLA that searches for possible sub-threshold lensed counterparts to known superthreshold gravitational waves. TESLA also participated in the LVK collaboration-wide effort to search for lensing signatures in gravitational waves from O3 data. Results from the analysis of the first-half of O3 data [23], the full O3 data [32], and the follow-up analysis [203] are included and discussed in Chapters 8, 9 and 10. In Chapters 11

¹ Under the strong lensing hypothesis, strongly-lensed gravitational waves should have the same intrinsic parameters. While the inferred sky location of the gravitational-wave sources of the lensed signals should differ, the deviation is usually of the order of $O(1'')$. In contrast, the source localization for gravitational waves has an uncertainty of the order of $O(1^\circ)$. Therefore, we assume that strongly-lensed gravitational waves would have essentially the same sky location.

and 12, I will discuss the latest improvements made to the TESLA pipeline (currently known as the TESLA-X pipeline) that boost the search sensitivity for possible sub-threshold lensed counterparts to superthreshold gravitational waves.

APPENDIX

6.A Gravitational lensing in Newtonian Mechanics

Escape velocity

In Newtonian Mechanics, when an object is being projected from the surface of the Earth, depending on the initial velocity, the object may

1. return to the Earth's surface,
2. move around the Earth in a bounded orbit, or
3. leave the Earth forever through an unbound orbit.




Figure 6.A.1: (Left) Objects projected from the Earth's surface can have very different trajectories depending on the initial conditions. (Right) Some objects in the universe may have accumulated sufficiently high masses. Gravity is so strong that not even light can escape.

For an object to escape the influence of gravity of its host, we can compute the escape velocity v_{esc} by considering energy conservation: The initial mechanical energy of the object must be conserved as it escapes from its host. Suppose an object of mass m is projected with an initial velocity of v from the surface of a host with radius R and mass M . If the object manages to escape from its host completely, the object should have zero gravitational potential energy, and the minimum final velocity of

the object is zero. Mathematically, we can write

$$\frac{1}{2}mv^2 - \frac{GMm}{R} = 0 \quad (6.54)$$

$$v = v_{\text{esc}} = \sqrt{\frac{2GM}{R}}, \quad (6.55)$$

where G is the gravitational constant. This is known as the escape velocity v_{esc} .

We can evaluate the radius of the host when the escape velocity becomes or exceeds the speed of light c :

$$c = \sqrt{\frac{2GM}{R}} \quad (6.56)$$

$$R = \frac{2GM}{c^2}. \quad (6.57)$$

Such radius is known as the Schwarzschild radius. Nothing, including light, can escape if a mass M is compressed to a radius smaller or equal to R . In the 18th Century, John Michell was the first person to theorize the existence of these “compact” objects, which he named “dark stars” [252].

The above discussion hints at the point that gravity affects not only masses but also light. When objects pass by masses, their travelling paths will get deflected due to gravity. The same goes for light, which will be bent as it passes by massive intervening objects. When these deflected light rays reach the Earth, they appear to originate from positions different from the source position. In this sense, the massive object that “supplies” the gravity for light bending becomes a “gravitational lens”. Figure 6.A.2 shows an illustrative example of gravitational lensing of a star-forming galaxy by a gravitational lens formed by a galaxy cluster. The images formed from gravitational lensing can also suffer a magnification / de-magnification in terms of brightness (amplitude) and size, and their shapes may also be deformed.

6.B Derivation of the deflection angle in Newtonian Mechanics

In this section, I present four different methods to derive the Newtonian deflection angle.

Method 1: Dimensional analysis

We note that the deflection angle $\hat{\alpha}$ is a dimensionless quantity. It is possible to construct this dimensionless physical quantity with known related parameters. We have earlier derived the Schwarzschild radius $R_S = \frac{2GM}{c^2}$, which has a unit of length. This is a good starting point because it relates the gravitational “strength” of the lens




Figure 6.A.2: Light coming from a star-forming galaxy passes through a galaxy cluster that acts as a gravitational lens to deflect the light rays. Observers on the other side of the lens will think that the deflected light comes from images at positions different from the source position. The images may be magnified / de-magnified in terms of brightness and size, and may suffer from shape deformation.

to its density (mass M and radius R). To create a dimensionless quantity, we must divide the Schwarzschild radius with another parameter with a unit of length. A convenient parameter would be the impact parameter ξ . Putting everything together, we get a dimensionless expression

$$\hat{\alpha} = \frac{2GM}{c^2\xi}, \quad (6.58)$$

which turns out to be exactly the Newtonian deflection angle. We can make sense of the constructed quantity:

1. The deflection angle is directly proportional to the mass of the lens M . That is, the more The larger the lens is, the larger the deflection angle will be, i.e. the more significant the gravitational lensing effect will be.
2. The deflection angle is inversely proportional to the impact parameter ξ . That is, the closer the source is to the principal axis, the greater the deflection angle will be. This makes perfect sense because gravitational potential is inversely proportional to the distance to the source of gravity. The closer it is to the source of gravity, the stronger the gravitational force will be, and hence, the greater the effect of gravitational lensing will be.

Method 2: A hand-waving approach




Figure 6.B.1: An astronaut is riding in a rocket originally moving at speed v horizontally to the right. A blue giant of mass M nearby acts as the gravitational lens. The *half*-deflection angle is $\hat{\alpha}$, and the distance d is the impact parameter. Inside the rocket, the astronaut shoots a light beam horizontally (in red).

Figure 6.B.1 shows an astronaut riding a rocket moving horizontally to the right with an initial speed v . A blue giant of mass M nearby acts as the gravitational lens that changes the rocket's trajectory. The *half*-deflection angle is $\hat{\alpha}$ ², and the impact parameter is d . Inside the rocket, the astronaut shines a light beam horizontally (in red) to the right. If we assume the change in distance between the blue giant and the rocket $d\vec{r}$ is small, the magnitude of the acceleration due to gravity $|\vec{a}| = \left| \frac{d^2\vec{r}}{dt^2} \right|$ can be regarded as approximately constant throughout the whole deflection process. Under this assumption, the acceleration can be approximated as

$$a \approx \frac{GM}{d^2}. \quad (6.59)$$

Since the rocket, astronaut, and the light pulse all experience the same deflection due to gravity. The time t that the light pulse experiences deflection due to gravity is roughly the same order as it takes to travel a distance d (the impact parameter) in its original speed (i.e. the speed of light). Hence, we can approximate t as

$$t \approx \frac{d}{c}. \quad (6.60)$$

² The deflection angle is halved because we have only considered half the deflection process.

As the acceleration is assumed to be roughly constant throughout the whole deflection, the final vertical speed of the light pulse is approximately

$$v_{\perp} = u + at = 0 + \frac{GM}{d^2} \times \frac{d}{c} = \frac{GM}{cd}, \quad (6.61)$$

where u is the initial vertical speed of the light pulse. Assuming that the half-deflection angle $\hat{\alpha}$ is small, and considering the horizontal component $v = c$ and vertical component v_{\perp} of the speed of the light pulse, we get

$$\tan \alpha \approx \hat{\alpha} = \frac{v_{\perp}}{v} = \frac{GM}{c^2 d}, \quad (6.62)$$

which is the same expression we got earlier (apart from the prefactor 2 because we have only considered half the deflection process).

Method 3: Integration




Figure 6.B.2: Here, we show a rocket (the rectangle box) moving horizontally along the positive x -axis initially. A blue giant of mass M (in blue) is placed at the origin. \vec{r} denotes the vector pointing from the blue giant to the rocket, and θ is the angle between \vec{r} and the positive y -axis. $\hat{\alpha}$ is the deflection angle, and d is the impact parameter. A light beam is shot horizontally inside the rocket (not shown).

Figure 6.B.2 shows a rocket (the rectangle box) moving horizontally initially along the positive x -axis. A blue giant of mass M (in blue) is placed at the origin. \vec{r} denotes the vector pointing from the blue giant to the rocket, and θ is the angle between \vec{r} and the positive y -axis. $\hat{\alpha}$ is the deflection angle, and d is the impact parameter. A light beam is shot horizontally inside the rocket (not shown).

Assuming the deflection is very small such that the y -component of the vector \vec{r} is approximately constant, i.e. it is approximately equal to the impact parameter d , we can express r as

$$r \approx \sqrt{x^2 + d^2}, \quad (6.63)$$

where x is the x -coordinate of the rocket. Neglecting the horizontal component of the gravitational acceleration on the light beam by the blue giant ³, then the acceleration due to gravity experienced by the light beam at any time t is given by

$$\vec{a} = -\frac{GM}{r^2}\hat{r} \approx -\frac{GM}{r^2}\cos\theta\hat{y}. \quad (6.64)$$

The magnitude of the acceleration can be further expressed as

$$a = \frac{GM}{r^2}\cos\theta = \frac{GM}{r^2}\frac{d}{r} = \frac{GMd}{(x^2 + d^2)^{3/2}}. \quad (6.65)$$

Then, the total increase in speed of the light beam along the y -direction is given by

$$\Delta v_{\perp} = \int_{t_0}^{t_{\infty}} a dt. \quad (6.66)$$

We make a change of variable with

$$dt = \frac{dx}{c}, \quad (6.67)$$

where dt is the infinitesimal time of travel of the light beam along the x direction. Since we assume the horizontal component of the acceleration due to gravity is negligible, the speed of the light beam along the x -direction remains c . Then, we get

$$\Delta v_{\perp} = \int_{-\infty}^{\infty} a \frac{dx}{c} = \frac{GMd}{c} \int_{-\infty}^{\infty} \frac{1}{(x^2 + d^2)^{3/2}} dx. \quad (6.68)$$

Doing a change of variable $x = d \tan \theta$, we get $dx = d \sec^2 \theta d\theta$, and hence we have

$$\Delta v_{\perp} = \frac{GMd}{c} \int_{-\pi/2}^{\pi/2} \frac{1}{d^3 \sec^3 \theta} d \sec^2 \theta d\theta \quad (6.69)$$

$$= \frac{GM}{cd} \int_{-\pi/2}^{\pi/2} \frac{d\theta}{\sec \theta} \quad (6.70)$$

$$= \frac{GM}{cd} \int_{-\pi/2}^{\pi/2} \cos \theta d\theta = \frac{2GM}{cd}. \quad (6.71)$$

³ We neglect that because we are only interested in the deflection along the vertical direction.

Assuming that the deflection angle is small, we have

$$\tan \alpha \approx \alpha = \frac{\Delta v_{\perp}}{c} = \frac{2GM}{c^2 d}, \quad (6.72)$$

which is the same expression we obtained earlier ⁴.

Method 4: Standard Newtonian Mechanics

Image




Figure 6.B.3: A gravitational lens of mass M is placed at the origin. The x -axis is aligned such that the positive- x direction is along the direction of the un-deflected light ray from the source. The distance from the lens to the closest approach of the light ray is ξ , at a polar angle θ_0 . The position of the light ray at any time t is represented by the position vector \vec{r} . The half-deflection angle is $\delta\theta$, so the total deflection angle is $2\delta\theta$ (deflection angle is defined as the angle between the image and the source).

Figure 6.B.3 shows a gravitational lens of mass M at the origin. The x -axis is aligned such that the positive- x direction is along the direction of the undeflected light ray from the source. The distance from the lens to the closest approach of the light ray is ξ , at a polar angle θ_0 . The position of the light ray at any time t is represented by the position vector \vec{r} . The half-deflection angle is $\delta\theta$, so the total deflection angle is $2\delta\theta$ (deflection angle is defined as the angle between the image and the source).

⁴ Note that this time, we get the correct prefactor 2 because the integral ranges from $-\infty$ to ∞ . Technically, we are integrating over the whole deflection process.

Assuming photons have mass, we can write the equations of motion for photons as

$$m \frac{d^2 \vec{r}}{dt^2} = -\frac{GMm}{r^2} \hat{r} \quad (6.73)$$

$$\frac{d^2 \vec{r}}{dt^2} = -\frac{GM}{r^2} \hat{r}, \quad (6.74)$$

where we have removed photons “hypothetical” mass m . This equation of motion is, in general a three-dimensional equation (in spherical coordinates, \hat{r} , $\hat{\theta}$, $\hat{\phi}$). However, we can construct the diagram such that all the motions are constrained on a plane, which reduces this to a two-dimensional equation (in polar coordinates, \hat{r} and $\hat{\theta}$). Recall that $\vec{r} = r\hat{r}$, $\dot{\vec{r}} = \dot{r}\hat{r} + r\dot{\theta}\hat{\theta}$, and $\ddot{\vec{r}} = (\ddot{r} - r\dot{\theta}^2)\hat{r} + (r\ddot{\theta} - 2\dot{r}\dot{\theta})\hat{\theta}$. We can hence write the above equation of motion as

$$\left(\ddot{r} - r\dot{\theta}^2 + \frac{GM}{r^2} \right) \hat{r} + (r\ddot{\theta} - 2\dot{r}\dot{\theta}) \hat{\theta} = 0. \quad (6.75)$$

Focusing on the θ component of the equation of motion, we have

$$r\ddot{\theta} - 2\dot{r}\dot{\theta} = 0 \quad (6.76)$$

$$r^2\ddot{\theta} - 2r\dot{r}\dot{\theta} = 0 \quad (6.77)$$

$$\frac{d}{dt} [r^2\dot{\theta}] = 0 \quad (6.78)$$

$$r^2\dot{\theta} = L_z = \text{Constant}, \quad (6.79)$$

which leads to a “conserved” quantity $L_z = r^2\dot{\theta}$, corresponding to the “angular momentum”.

The r -component of the equation of motion can be written as

$$\ddot{r} - \frac{L_z^2}{r^3} + \frac{GM}{r^2} = 0. \quad (6.80)$$

To solve the equation, we make a change of variable $r(t) \rightarrow r[\theta(t)]$, i.e. we make r become a function of θ instead of time t . Then, we have

$$\dot{r} = \frac{dr}{d\theta} \frac{d\theta}{dt} = r'\dot{\theta} = r' \frac{L_z}{r^2}, \quad (6.81)$$

where \dot{r} and r' represent the derivative with respect to time t and θ respectively. Similarly,

$$\ddot{r} = L_z \frac{d}{dt} \left[r' \frac{1}{r^2} \right] = L_z \frac{d}{d\theta} \left[r' \frac{1}{r^2} \right] \times \frac{d\theta}{dt} = \frac{L_z^2}{r^2} \left[\frac{r''}{r^2} - \frac{2r'^2}{r^3} \right] = L_z^2 \left[\frac{r''}{r^4} - \frac{2r'^2}{r^5} \right]. \quad (6.82)$$

The equation of motion can then be written as

$$L_z^2 \left[\frac{r''}{r^4} - \frac{2r'^2}{r^5} \right] + \frac{GM}{r^2} = 0. \quad (6.83)$$

We make another change of variable $u = \frac{1}{r}$. This gives us

$$r = \frac{1}{u} \rightarrow r' = -\frac{1}{u^2}u' \rightarrow r'' = -\frac{u''}{u^2} + \frac{2u'^2}{u^3}, \quad (6.84)$$

and the equation of motion becomes

$$L_z^2 (u^2 u'' + u^3) = GM u^2 \quad (6.85)$$

$$u^2 [L_z^2 (u'' + u) - GM] = 0. \quad (6.86)$$

This leads to two possible solutions:

1. $u = 0$ and hence $r \rightarrow \infty$ (the light is infinitely away from the lens, such that no lensing occurs), and
2. $L_z^2 (u'' + u) - GM = 0$.

The first solution is the trivial solution (no lensing has occurred). Hence, we will focus on the second solution. The second solution yields an inhomogeneous second-order differential equation

$$u'' + u = \frac{GM}{L_z^2}. \quad (6.87)$$

The homogenous part of the equation

$$u'' + u = 0 \quad (6.88)$$

yields the homogenous solution $u_h = A \cos(\theta - \phi)$, where A and ϕ are constants to be determined by initial conditions. The particular solution is simply the constant on the right side of the equation, i.e. $u_p = \frac{GM}{L_z^2}$. Hence, the general solution to the differential equation is given by

$$u = u_h + u_p = A \cos(\theta - \phi) + \frac{GM}{L_z^2}. \quad (6.89)$$

In terms of r , the solution is

$$\frac{1}{r(\theta)} = A \cos(\theta - \phi) + \frac{GM}{L_z^2}. \quad (6.90)$$

We observe that r is minimized when the right side is maximized, which, by constraining A to be positive, happens when $\cos(\theta - \phi) = 1$ (i.e. $\theta - \phi = 0$). The minimum r occurs at the closest approach $\theta = \theta_0$. Hence, we immediately get $\phi = \theta_0$, and the solution becomes

$$\frac{1}{r(\theta)} = A \cos(\theta - \theta_0) + \frac{GM}{L_z^2}. \quad (6.91)$$

Now, we consider the conservation of angular momentum $L_z = r^2 \dot{\theta}$. When the light is very far away from the lens $\theta \rightarrow \theta_\infty$, the angular momentum is roughly equal to $L_z = \vec{r} \times \vec{v} \approx \xi c$, since the light is moving along the direction of x . We can hence write the solution for r as

$$\frac{1}{r(\theta)} = A \cos(\theta - \theta_0) + \frac{GM}{\xi^2 c^2}. \quad (6.92)$$

We now impose initial conditions: (1) When $\theta \rightarrow \pi$, $r \rightarrow \infty$, so we get

$$0 = A \cos(\pi - \theta_0) + \frac{GM}{\xi^2 c^2} = -A \cos \theta_0 + \frac{GM}{\xi^2 c^2} \quad (6.93)$$

$$A \cos \theta_0 = \frac{GM}{\xi^2 c^2} \quad (6.94)$$

$$A^2 \cos^2 \theta_0 = \frac{G^2 M^2}{\xi^4 c^4}. \quad (6.95)$$

(2) If we differentiate the equation of motion with respect to time, we get

$$-\frac{\dot{r}}{r^2} = -A \dot{\theta} \sin(\theta - \theta_0) \quad (6.96)$$

$$A = \frac{\dot{r}}{\underbrace{r^2 \dot{\theta}}_{L_z} \sin(\theta - \theta_0)} = \frac{\dot{r}}{\xi c \sin(\theta - \theta_0)}. \quad (6.97)$$

Initially at $\theta = \pi$, we have $\dot{\theta} \approx 0$, and recall that $v = c = \sqrt{\dot{r}^2 + r^2 \dot{\theta}^2} \approx \dot{r}$ in polar coordinates, so we get

$$A = \frac{c}{\xi c \sin(\pi - \theta_0)} = \frac{1}{\xi \sin \theta_0} \quad (6.98)$$

$$\sin \theta_0 = \frac{1}{A \xi} \quad (6.99)$$

$$\cos^2 \theta_0 = 1 - \frac{1}{A^2 \xi^2}. \quad (6.100)$$

Together, we get

$$A^2 \left(1 - \frac{1}{A^2 \xi^2} \right) = \frac{G^2 M^2}{\xi^4 c^4} \quad (6.101)$$

$$A^2 = \frac{1}{\xi^2} + \frac{G^2 M^2}{\xi^4 c^4}, \quad (6.102)$$

and hence

$$\cos^2 \theta_0 = 1 - \frac{1}{\xi^2 \left(\frac{1}{\xi^2} + \frac{G^2 M^2}{\xi^4 c^4} \right)} = 1 - \frac{1}{1 + \frac{G^2 M^2}{\xi^2 c^4}} \approx 1 - \left[1 - \frac{G^2 M^2}{\xi^2 c^4} \right] = \frac{G^2 M^2}{\xi^2 c^4} \quad (6.103)$$

$$\cos \theta_0 = \frac{GM}{\xi c^2}. \quad (6.104)$$

Now note that if $\theta_0 = \frac{\pi}{2}$, it results in no deflection. Therefore, if we write $\theta_0 = \frac{\pi}{2} - \epsilon$, where ϵ is a small number, then ϵ will define the “strength” of the lensing effect acting on the light beam. Then, we can expand $\cos \theta_0$ as

$$\cos \theta_0 = \cos \left(\frac{\pi}{2} - \epsilon \right) \approx \epsilon = \frac{GM}{\xi c^2}, \quad (6.105)$$

and hence

$$\theta_0 \approx \frac{\pi}{2} - \frac{GM}{\xi c^2}. \quad (6.106)$$

Finally, consider the geometric construction in Figure 6.B.4. In the closed quadrilateral, by the angle sum of polygon, we have

$$\frac{\pi}{2} \times 2 + \left(\frac{\pi}{2} - \delta\theta \right) + (\pi - \theta_0) = 2\pi \quad (6.107)$$

$$\delta\theta = \frac{\pi}{2} - \left(\frac{\pi}{2} - \frac{GM}{\xi c^2} \right) = \frac{GM}{\xi c^2}, \quad (6.108)$$

and hence the total deflection angle is, once again,

$$\hat{\alpha} = 2\delta\theta = \frac{2GM}{\xi c^2}. \quad (6.109)$$

6.C Relativistic deflection angle

In this section, we will explain how one can derive the relativistic deflection angle. Unless otherwise stated, we set the speed of light $c = 1$ for simplicity. We consider a simple point-mass lens massive enough to be considered stationary. Around a point mass lens of mass M , spacetime can be described by the Schwarzschild metric, i.e.

$$ds^2 = - \left(1 - \frac{2GM}{r} \right) dt^2 + \left(1 - \frac{2GM}{r} \right)^{-1} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \quad (6.110)$$




Figure 6.B.4: Additional geometric construction on 6.B.3.

in spherical coordinates. Note that away from the point mass lens (i.e. $r \rightarrow \infty$), the Schwarzschild metric reduces to the flat Minkowski metric, i.e.

$$ds^2 \rightarrow -dt^2 + dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2. \quad (6.111)$$

The metric tensor $g_{\mu\nu}$ for the Schwarzschild metric can be written as

$$g_{\mu\nu} = \text{diag} \left[-\left(1 - \frac{2GM}{r}\right), \left(1 - \frac{2GM}{r}\right)^{-1}, r^2, r^2 \sin^2 \theta \right]. \quad (6.112)$$

The geodesic equation is given by

$$\frac{d}{d\tau} \left(g_{\mu\nu} \frac{dx^\nu}{d\tau} \right) - \frac{1}{2} g_{\alpha\beta,\mu} \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0, \quad (6.113)$$

where $g_{\alpha\beta,\mu} = \partial_\mu g_{\alpha\beta}$, τ is the proper time, and x^μ is the 4-position vector. Considering the $\mu = t$ -component of the geodesic equation, we have

$$\frac{d}{d\tau} \left(g_{t\nu} \frac{dx^\nu}{d\tau} \right) - \frac{1}{2} g_{\alpha\beta,t} \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0. \quad (6.114)$$

Since $g_{t\nu}$ is non-zero only when $\nu = t$, and none of the components of $g_{\mu\nu}$ is explicitly dependent on t , we have $g_{\alpha\beta,t} = 0$, and

$$\frac{d}{d\tau} \left[g_{it} \frac{dt}{d\tau} \right] = 0 \quad (6.115)$$

$$-g_{it} \frac{dt}{d\tau} = e = \text{Constant}. \quad (6.116)$$

A negative sign is deliberately multiplied to the last line of the equation. Note that $\frac{dt}{d\tau}$ is the t -component of the 4-velocity v^t . If we multiply the mass of the corresponding particle m to the expression, we get the t -component p^t of the 4-momentum. As $r \rightarrow \infty$, $g_{tt} \rightarrow -1$ and $me = p^t$, corresponding to the energy of the particle. Therefore, we can interpret e as the energy per unit mass of the particle.

Considering the $\mu = \phi$ component of the geodesic equation, we get

$$\frac{d}{d\tau} \left(g_{\phi\nu} \frac{dx^\nu}{d\tau} \right) - \frac{1}{2} g_{\alpha\beta,\phi} \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0. \quad (6.117)$$

Note that $g_{\phi\nu}$ is non-zero only when $\nu = \phi$, and $g_{\alpha\beta,\phi} = 0$. We have

$$\frac{d}{d\tau} \left(g_{\phi\phi} \frac{dx^\phi}{d\tau} \right) = 0 \quad (6.118)$$

$$g_{\phi\phi} \frac{dx^\phi}{d\tau} = r^2 \sin^2 \theta \frac{dx^\phi}{d\tau} = l = \text{Constant}. \quad (6.119)$$

Assuming the geodesic lies on the equatorial plane, i.e. $\theta = \frac{\pi}{2}$, we get

$$l = r^2 \frac{dx^\phi}{d\tau} = r^2 \dot{\phi}. \quad (6.120)$$

Multiplying the mass of the particle m to both sides, we get

$$ml = mr^2 \dot{\phi}, \quad (6.121)$$

which we can immediately identify the right side as the angular momentum of the particle. Therefore, we can interpret l as the particle's angular momentum per unit mass.

Note that the above discussion only applies to particles with mass. For massless particles like photons, the proper time τ is ill-defined (along any photon's worldline, $\tau = 0$). We can eliminate the proper time τ by creating a new conserved quantity $b = \frac{l}{e}$ using the previously found conserved quantities. This new quantity

$$b = \frac{l}{e} = \frac{r^2 \frac{d\phi}{d\tau}}{\left(1 - \frac{2GM}{r}\right) \frac{dt}{d\tau}} = \frac{r^3}{r - 2GM} \frac{d\phi}{dt} \quad (6.122)$$

$$\frac{d\phi}{dt} = \frac{b}{r^2} \left(1 - \frac{2GM}{r}\right). \quad (6.123)$$

is the impact parameter. Note that this new quantity b does not depend on τ , hence it applies to massless particles like photons.

For massless particles like photons, the proper time is always $d\tau = 0$, so by comparing the spacetime interval measured by an observer co-moving with a photon (i.e. the proper time $-d\tau = 0$), and that measured by another arbitrary observer, we get

$$0 = -\left(1 - \frac{2GM}{r}\right) dt^2 + \left(1 - \frac{2GM}{r}\right)^{-1} dr^2 + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2. \quad (6.124)$$

Constraining the motion on the equatorial plane (i.e. $\theta = \frac{\pi}{2}$ and $d\theta = 0$), we get

$$0 = -\left(1 - \frac{2GM}{r}\right) dt^2 + \left(1 - \frac{2GM}{r}\right)^{-1} dr^2 + r^2 d\phi^2 \quad (6.125)$$

$$0 = -\left(1 - \frac{2GM}{r}\right) + \left(1 - \frac{2GM}{r}\right)^{-1} \left(\frac{dr}{dt}\right)^2 + r^2 \left(\frac{d\phi}{dt}\right)^2. \quad (6.126)$$

Using the quantity b we obtained earlier, we have

$$0 = -\left(1 - \frac{2GM}{r}\right) + \left(1 - \frac{2GM}{r}\right)^{-1} \left(\frac{dr}{dt}\right)^2 + r^2 \times \frac{b^2}{r^4} \left(1 - \frac{2GM}{r}\right)^2 \quad (6.127)$$

$$\left(\frac{dr}{dt}\right)^2 = \left(1 - \frac{2GM}{r}\right)^2 - \frac{b^2}{r^2} \left(1 - \frac{2GM}{r}\right)^3 \quad (6.128)$$

$$\frac{dr}{dt} = \left(1 - \frac{2GM}{r}\right) \sqrt{1 - \frac{b^2}{r^2} \left(1 - \frac{2GM}{r}\right)}. \quad (6.129)$$

Making a change of variable $u(\phi) = \frac{1}{r(\phi)}$ gives

$$\frac{dr}{dt} = \frac{dr}{du} \frac{du}{d\phi} \frac{d\phi}{dt} = -\frac{1}{u^2} \times u' \times \frac{b}{r^2} \left(1 - \frac{2GM}{r}\right) = -u' b (1 - 2GMu), \quad (6.130)$$

where $u' = \frac{du}{d\phi}$. Then, the previous equation can be rewritten as

$$-u' b (1 - 2GMu) = (1 - 2GMu) \sqrt{1 - b^2 u^2 (1 - 2GMu)} \quad (6.131)$$

$$(1 - 2GMu) \left[u' b + \sqrt{1 - b^2 u^2 (1 - 2GMu)} \right] = 0, \quad (6.132)$$

which leads to

$$u' b + \sqrt{1 - b^2 u^2 (1 - 2GMu)} = 0 \quad (6.133)$$

$$u'^2 b^2 = 1 - b^2 u^2 (1 - 2GMu) \quad (6.134)$$

$$u'^2 + u^2 = \frac{1}{b^2} + 2GMu^3. \quad (6.135)$$

Differentiating both sides with respect to ϕ again, we get

$$2u'u'' + 2uu' = 6GMu^2u' \quad (6.136)$$

$$u'(u'' + u - 3GMu^2) = 0. \quad (6.137)$$

This leads to either $u' = 0$ (trivial solution, corresponding to the scenario where no lensing occurs), or the differential equation

$$u'' + u = 3GMu^2. \quad (6.138)$$

To solve this differential equation, we first consider the case where $M = 0$, i.e. when there is no lens. Then, the differential equation becomes

$$u'' = -u \quad (6.139)$$

and the general solution is simply

$$\tilde{u} = A \cos \phi + B \sin \phi, \quad (6.140)$$

where A, B are constants to be determined by initial conditions. If we assume $u = 0$ (i.e. $r \rightarrow \infty$) at $\phi = \pi$, we get $A = 0$ and $\tilde{u} = B \sin \phi$. Note that u is maximized when $\phi = \frac{\pi}{2}$, which corresponds to the minimum of r , i.e. the closest approach u_c from the gravitational lens to the light beam. With such, we immediately get $A = u_c$ and hence $\tilde{u} = u_c \sin \phi$.

With perturbation theory, we can now solve the original equation (i.e. $M \neq 0$). We assume the solution can be written as

$$u = u_c \sin \phi + u_c z(\phi), \quad (6.141)$$

where $z(\phi)$ is a small perturbation. Substituting this into the original equation, we get

$$\frac{d^2}{d\phi^2} [u_c \sin \phi + u_c z(\phi)] + [u_c \sin \phi + u_c z(\phi)] = 3GM [u_c \sin \phi + u_c z(\phi)]^2 \quad (6.142)$$

$$-u_c \sin \phi + u_c z'' + u_c \sin \phi + u_c z = 3GM [u_c^2 \sin^2 \phi + 2u_c^2 z \sin \phi + u_c^2 z^2] \quad (6.143)$$

$$z'' + z \approx 3GM(u_c \sin^2 \phi) \quad (6.144)$$

where we have dropped the final two terms on the right side because they are sufficiently small. This means

$$z'' + z = \frac{3GMu_c}{2} (1 - \cos 2\phi). \quad (6.145)$$

This is, again, an inhomogeneous differential equation. To solve for the particular solution, we guess the solution to be

$$z_p = E + F \sin 2\phi + H \cos 2\phi \quad (6.146)$$

where E, F, H are constants. Substituting this into the equation, we get

$$[-4F \sin 2\phi - 4H \cos 2\phi] + [E + F \sin 2\phi + H \cos 2\phi] = \frac{3GMu_c}{2} (1 - \cos 2\phi) \quad (6.147)$$

$$-3F \sin 2\phi - 3H \cos 2\phi + E = \frac{3GMu_c}{2} (1 - \cos 2\phi). \quad (6.148)$$

This gives $F = 0$, $E = \frac{3GMu_c}{2}$ and $H = \frac{GMu_c}{2}$. Then, the general solution for z can be written as

$$z = \frac{3GMu_c}{2} + \frac{GMu_c}{2} \cos 2\phi, \quad (6.149)$$

and that for u can be written as

$$u = u_c \sin \phi + \frac{3GMu_c^2}{2} + \frac{GMu_c^2}{2} \cos 2\phi = u_c \left[\sin \phi + \frac{3GMu_c}{2} \left(1 + \frac{1}{3} \cos 2\phi \right) \right]. \quad (6.150)$$

Recall that $u = 0$ when $\phi = 0$ and π for $M = 0$. However, this is no longer the case when $M \neq 0$. Instead, $u = 0$ for some small angle ϕ_0 (ϕ_0 is small because we assume a small perturbation). Under the small angle approximation, we get $\sin \phi \approx \phi$ and $\cos 2\phi \approx 1$, and hence

$$0 \approx u_c \left[\phi_0 + \frac{3GMu_c}{2} \left(1 + \frac{1}{3} \right) \right] \quad (6.151)$$

$$\phi_0 = 2GMu_c = \frac{2GM}{r_c}, \quad (6.152)$$

where $r_c = \frac{1}{u_c}$ is the distance from the lens to the light beam at the closest approach. Since we are only considering half the deflection process, the total deflection angle is hence

$$\hat{\alpha} = \frac{4GM}{r_c}, \quad (6.153)$$

which is twice the Newtonian value.

6.D The lens equation from geometric constructions

In this section, we will derive the lens equation with the thin lens approximation.

The 1D lens equation

We first consider the “simplest” lens model - When there is no lens. Figure 6.D.1 illustrates the scenario. Although no lensing is occurring, we can already define two important parameters, namely the source position β and source distance D_S , for later discussion. Note that all distances in this section are assumed to be cosmological distances.




Figure 6.D.1: The “simplest” lens model is the case when no gravitational lens exists. In this scenario, however, we can already define two important quantities, namely the source position β and source distance D_S , which will be useful later.

We now insert a massive lens into the picture. We assume that the lens is massive enough so it is approximately a stationary point mass. Figure 6.D.2 illustrates the updated scenario. In this case, we need to define several new parameters, namely (1) the image position θ , (2) the lens-source distance D_{LS} , (3) the lens distance D_L , (4) the impact parameter ξ , which is the distance to the closest approach between the light ray and the lens, and (5) the deflection angle $\hat{\alpha}$, which is the angle between the image and the undeflected light ray.

Recall that the deflection angle is given by

$$\hat{\alpha} = \frac{4GM}{\xi}. \quad (6.154)$$

Assuming small angles, the “height” of the image is then given by $D_S\theta$. The same “height” can also be expressed in terms of D_S , D_{LS} , β and $\hat{\alpha}$ as $D_S\beta + D_{LS}\hat{\alpha}$. The impact parameter can be expressed as $\xi = D_L\theta$. Hence, we have

$$D_S\theta = D_S\beta + D_{LS}\hat{\alpha} = D_S\beta + \frac{4GMD_{LS}}{D_L\theta} \quad (6.155)$$

$$\theta^2 - \beta\theta - \frac{4GMD_{LS}}{D_LD_S} = 0. \quad (6.156)$$




Figure 6.D.2: A massive point-mass lens enters the picture. Because of gravitational lensing, the light ray from the source is bent, and a new image is formed. We now need to define several new parameters: (1) The image position θ , (2) The lens-source distance D_{LS} , (3) The lens distance D_L , (4) The impact parameter ξ (the distance of closest approach between the light ray and the lens), and (5) The deflection angle $\hat{\alpha}$ (the angle between the un-deflected light ray and the image).

This quadratic equation for θ is known as the “lens equation”. If we set $\beta = 0$ (i.e. the source is directly behind the lens), we get

$$\theta^2 = \frac{4GM D_{LS}}{D_L D_S} \quad (6.157)$$

$$\theta = \theta_E = \pm \sqrt{\frac{4GM D_{LS}}{D_L D_S}}. \quad (6.158)$$

θ_E is known as the “Einstein radius”. In this case, “infinitely many” images are formed around the point mass lens with the image position being θ_E , forming the “Einstein ring”.

We can then write the lens equation in terms of θ_E as

$$\theta^2 - \beta\theta - \theta_E^2 = 0 \quad (6.159)$$

$$\theta = \frac{\beta}{2} \pm \frac{1}{2} \sqrt{\beta^2 + 4\theta_E^2} = \theta_{\pm}, \quad (6.160)$$

which gives the image positions as θ_{\pm} . Another way to visualize the result is by using the graphical method. The lens equation can be written as

$$\theta - \beta = \frac{\theta_E^2}{\theta}. \quad (6.161)$$

The problem is now transformed into one asking for the number of intersections between the two curves $y = \theta - \beta$ and $y = \frac{\theta_E^2}{\theta}$. As Figure 6.D.3 demonstrates, the point mass lens model always admits two solutions.




Figure 6.D.3: The problem of solving the image positions is transformed into a graphical problem, asking for the number of intersection points between two curves $y = \theta - \beta$ and $y = \frac{\theta_E^2}{\theta}$. As we can see, the point mass lens model always has two images (two intersection points in red).

Lens equation in higher dimensions

We now want to extend our discussion to three-dimensional scenarios where the lens has sizes. For the following derivation, refer to Figure 6.2.2. All the angles are now promoted to 2D angles. The deflection angle is now a function of the impact parameter $\hat{\alpha} = \hat{\alpha}(\vec{\xi})$. Apart from this, the previous derivation for the 1D lens equation still holds as long as we promote the angles to two-dimensional, i.e.

$$\vec{\beta} + \frac{D_{\text{LS}}}{D_{\text{S}}} \hat{\alpha} = \vec{\theta}. \quad (6.162)$$

This is the 2D lens equation.

We can define the reduced deflection angle as

$$\vec{\alpha} = \frac{D_{\text{LS}}}{D_{\text{S}}} \hat{\alpha}. \quad (6.163)$$

Rearranging the terms, we get

$$D_S \vec{\alpha} = D_{LS} \hat{\alpha}. \quad (6.164)$$

Both sides roughly describe the distance between the source and the image, and $\vec{\alpha}$ is roughly the angle between the line joining the observer and the source, and that joining the observer and the image.

Since all angles are now two-dimensional, including the deflection angle $\hat{\alpha}$, we can write the deflection angle as a directional vector in terms of the impact parameter vector, i.e.

$$\hat{\alpha} = \frac{4GM}{\xi} \hat{\xi}. \quad (6.165)$$

Generalizing to a lens that is not at the origin (but at some position $\vec{\xi}'$), the deflection angle is

$$\hat{\alpha} = \frac{4GM}{|\vec{\xi} - \vec{\xi}'|^2} (\vec{\xi} - \vec{\xi}'). \quad (6.166)$$

Note that this applies only to a single point mass lens.




Figure 6.D.4: Deriving the deflection angle for a thin lens.

If the lens mass can be described by some continuous mass density function $\rho(\vec{r})$, the deflection angle can be written as a triple integral, i.e.

$$\hat{\alpha} = 4GM \int_V \frac{\rho(\vec{r}') d^3 r'}{|\vec{\xi} - \vec{\xi}'|^2} (\vec{\xi} - \vec{\xi}') = 4GM \int_{z_0}^{z_1} dr'_3 \int_S \frac{\rho(\vec{\xi}, r'_3) d^2 \xi}{|\vec{\xi} - \vec{\xi}'|^2} (\vec{\xi} - \vec{\xi}'), \quad (6.167)$$

where r_3 is the dimension perpendicular to the lens plane. Note that in principle we should replace the vector $\vec{\xi} - \vec{\xi}'$ by $\vec{\eta} - \vec{\eta}'$, where $\vec{\eta} = \vec{\xi} + \vec{r}_3$ includes the distance along the dimension r_3 . However, here we take the thin lens approximation (i.e. the distance along the dimension r_3 is negligible).

The integral along the r_3 dimension can be integrated easily since only ρ depends on r_3 . We can define the integral as a new quantity

$$\Sigma(\vec{\xi}) = \int \rho(\vec{\xi}, r'_3) dr'_3, \quad (6.168)$$

known as the surface mass density of the lens. Then, the deflection angle can be written as

$$\hat{\alpha} = 4G \int \Sigma \frac{\vec{\xi} - \vec{\xi}'}{|\vec{\xi} - \vec{\xi}'|^2} d^2\xi. \quad (6.169)$$

We have shown that $\vec{\xi} = D_L \vec{\theta}$. Similarly, we can write $\vec{\xi}' = D_L \vec{\theta}'$. Then, the deflection angle can be expressed as

$$\hat{\alpha} = 4G \int \Sigma \frac{D_L (\vec{\theta} - \vec{\theta}')}{D_L^2 |\vec{\theta} - \vec{\theta}'|^2} D_L^2 d^2\theta' = 4GD_L \int \Sigma \frac{\vec{\theta} - \vec{\theta}'}{|\vec{\theta} - \vec{\theta}'|^2} d^2\theta'. \quad (6.170)$$

The reduced deflection angle can also be expressed as

$$\vec{\alpha} = \frac{D_{LS}}{D_S} \hat{\alpha} = \frac{4GD_{LS}D_L}{D_S} \int \Sigma \frac{\vec{\theta} - \vec{\theta}'}{|\vec{\theta} - \vec{\theta}'|^2} d^2\theta'. \quad (6.171)$$

Note that the constants outside of the integral $\frac{4GD_{LS}D_L}{D_S}$ have a dimension of kg^{-1}m^2 (with appropriate insertion of the speed of light c), which is the inverse of the lens surface mass density. Therefore, we are prompted to define a “critical mass density” as

$$\Sigma_c = \frac{1}{4\pi G} \frac{D_S}{D_{LS}D_L}. \quad (6.172)$$

The constant π is inserted intentionally. Then, the reduced deflection angle can be written as

$$\vec{\alpha} = \frac{1}{\pi} \int \Sigma \frac{\vec{\theta} - \vec{\theta}'}{\Sigma_c |\vec{\theta} - \vec{\theta}'|^2} d^2\theta'. \quad (6.173)$$

We can further define “convergence” as

$$\kappa = \frac{\Sigma}{\Sigma_c}, \quad (6.174)$$

which allows us to further simplify the expression to

$$\vec{\alpha} = \frac{1}{\pi} \int \kappa \frac{\vec{\theta} - \vec{\theta}'}{|\vec{\theta} - \vec{\theta}'|^2} d^2\theta'. \quad (6.175)$$

From vector calculus, we know that

$$\nabla \ln |\vec{r} - \vec{r}'| = \frac{\vec{r} - \vec{r}'}{|\vec{r} - \vec{r}'|^2}, \quad (6.176)$$

so we can rewrite the reduced deflection angle as

$$\vec{\alpha} = \nabla \left[\frac{1}{\pi} \int \kappa \ln |\vec{\theta} - \vec{\theta}'| d^2\theta' \right]. \quad (6.177)$$

The term inside the square bracket is now solely a function of the image position θ , and all the lens information is well-contained in κ . We name the term the lens potential $\psi(\vec{\theta})$. Then, the reduced deflection angle can be written as

$$\vec{\alpha} = \nabla \psi(\vec{\theta}). \quad (6.178)$$

Many similarities between the lens equation derivations and electrostatic equations can be drawn. For instance, the electric field \vec{E} is minus the gradient of the electric potential V , i.e.

$$\vec{E} = -\nabla V. \quad (6.179)$$

For the case of lensing, the reduced deflection angle $\vec{\alpha}$ is the gradient of the lens potential $\psi(\vec{\theta})$, i.e.

$$\vec{\alpha} = \nabla \psi. \quad (6.180)$$

The total electric field by a distribution of charges is given by

$$\vec{E} = k_e \int \frac{\rho}{|\vec{r} - \vec{r}'|^3} (\vec{r} - \vec{r}') d^3r'. \quad (6.181)$$

For lensing, the reduced deflection angle due to a distribution of lens masses is given by

$$\vec{\alpha} = \frac{1}{\pi} \int \frac{\kappa}{|\vec{\theta} - \vec{\theta}'|^2} (\vec{\theta} - \vec{\theta}') d^2\theta'. \quad (6.182)$$

The Gauss's law for electrostatics is

$$\nabla \cdot \vec{E} = -\nabla^2 V = \frac{\rho}{\epsilon_0}. \quad (6.183)$$

The Gauss's law for lensing is

$$\nabla \cdot \vec{\alpha} = \nabla^2 \psi = 2\kappa. \quad (6.184)$$

The lens equations have almost the same structure as electrostatics formalism. Therefore, techniques for solving electrostatics problems can also be applied (with appropriate modifications) to solve the lens equation.

6.E Deriving the time delay components induced by gravitational lensing

We derive the lens equation by considering the total time delay induced by gravitational lensing. Gravitational lensing changes the path lengths of the waves traveling toward us from the source, causing a “geometric time delay”. There is also an additional “gravitational time delay” (also known as Shapiro time delay) that comes purely from the metric.

Geometric time delay

Let us first derive the geometric time delay due to the change in path length by gravitational lensing. We refer to Figure 6.E.1 for the derivation. By construction,




Figure 6.E.1: Deriving the geometric time delay.

we make $OA = OS$. Then, considering the isosceles triangle $\triangle OAS$, we have

$$\angle OAS = \angle OSA = \frac{\pi}{2} - \frac{\vec{\alpha}}{2}. \quad (6.185)$$

Next, we have $LI \approx LS \approx D_{LS}$, so by considering the isosceles triangle $\triangle LIS$, we have

$$\angle LIS = \angle LSI = \frac{\pi}{2} - \frac{\hat{\alpha}}{2}. \quad (6.186)$$

By considering the exterior angle of triangle $\triangle OLS$, we have

$$\angle OSL = \hat{\alpha} - \vec{\alpha}. \quad (6.187)$$

Then

$$\angle ASL = \frac{\pi}{2} - \frac{\vec{\alpha}}{2} - (\hat{\alpha} - \vec{\alpha}) = \frac{\pi}{2} + \frac{\vec{\alpha}}{2} - \hat{\alpha}. \quad (6.188)$$

With that, we get

$$\angle ASI = \frac{\pi}{2} - \frac{\hat{\alpha}}{2} - \left(\frac{\pi}{2} + \frac{\vec{\alpha}}{2} - \hat{\alpha} \right) = \frac{\hat{\alpha}}{2} - \frac{\vec{\alpha}}{2}. \quad (6.189)$$

Recall that the reduced deflection angle and the deflection angle are related by

$$\vec{\alpha} = \frac{D_{LS}}{D_S} \hat{\alpha}. \quad (6.190)$$

Therefore, we have

$$\angle ASI = \frac{D_S}{2D_{LS}} \vec{\alpha} - \frac{\vec{\alpha}}{2} = \frac{\vec{\alpha}}{2} \left(\frac{D_S}{D_{LS}} - 1 \right) = \frac{\vec{\alpha}}{2} \left(\frac{D_S - D_{LS}}{D_{LS}} \right) \approx \frac{\vec{\alpha}}{2} \frac{D_L}{D_{LS}}, \quad (6.191)$$

where in the last line we use the approximation $D_S - D_{LS} \approx D_L$. Now, $AS \approx D_S \vec{\alpha}$, so the additional path length is roughly

$$\Delta\lambda \approx D_S \vec{\alpha} \cdot \frac{\vec{\alpha}}{2} \frac{D_L}{D_{LS}} = \frac{D_S D_L}{2D_{LS}} |\vec{\alpha}|^2 = \frac{D_S D_L}{2D_{LS}} |\vec{\theta} - \vec{\beta}|^2. \quad (6.192)$$

We need to account for cosmological redshift near the lens, assuming that lensing happens only around the lens. Therefore, the geometric time delay is given by

$$\tau_{\text{geo}} = (1 + z_L) \Delta\lambda = (1 + z_L) \frac{D_S D_L}{2D_{LS}} |\vec{\theta} - \vec{\beta}|^2, \quad (6.193)$$

where z_L is the lens redshift.

The gravitational time delay (Shapiro time delay)

Another component for the gravitational lensing time delay comes purely from the metric, known as the gravitational time delay (Shapiro time delay). We will derive the Shapiro time delay in this subsection.

The metric of spacetime in which a lens with Newtonian potential Φ is present can be described by

$$ds^2 \approx -(1 + 2\Phi)dt^2 + (1 + 2\Phi)^{-1}(dx^2 + dy^2 + dz^2) \quad (6.194)$$

$$= -(1 + 2\Phi)dt^2 + (1 + 2\Phi)^{-1}dl^2. \quad (6.195)$$

For photons, the spacetime interval is always 0, so we have

$$0 = -(1 + 2\Phi)dt^2 + (1 + 2\Phi)^{-1}dl^2 \quad (6.196)$$

$$\frac{dl}{dt} = c' \approx 1 + 2\Phi. \quad (6.197)$$

Here, c' is the speed of the light near the lens. Note that $\Phi < 0$, meaning that light moves with a lower speed near the lens. Now, the total traveling time of light over its whole journey is given by

$$\int dt = t = \int \frac{dl}{1 + 2\Phi} \approx \int (1 - 2\Phi) dl = L - 2 \int \Phi dl, \quad (6.198)$$

where L is the time of travel if the speed of light is constant. So, the extra time required for light to travel the same path near a lens is given by

$$\tau_{\text{grav}} = -2(1 + z_L) \int \Phi dl, \quad (6.199)$$

where the factor $1 + z_L$ is inserted to account for cosmological redshift near the lens.

Chapter 7

THE TARGETED SUBTHRESHOLD SEARCH FOR STRONGLY LENSED GWS (TESLA)

Note: This Chapter is an adaptation of the short author publication:

Alvin K. Y. Li et al., “Targeted subthreshold search for strongly lensed gravitational-wave events”, *Phys. Rev. D* 107.12 (2023), doi: 10.1103/PhysRevD.107.123014

Alvin K. Y. Li is the first author of this paper. He is the main developer of the “TargetEd Subthreshold Lensing seArch” (TESLA) method to find possible sub-threshold lensed counterparts to known gravitational waves described in this manuscript and led the writing of the manuscript.

7.1 Introduction

At the time of writing, the LVK Collaboration published its first full-scale analysis to search for gravitational-lensing signatures of gravitational waves within data from the first half of LIGO/Virgo third observing run O3a [23]. They conclude that no compelling evidence was found for gravitational lensing to take place within O3a. In the paper, they consider the possibility that strong lensing produces multiple gravitational waves from the same sources. In one scenario some images are magnified and hence become identifiable as detections, and the rest are demagnified and thus are buried within the noise background. Through two independent search methods, they search for the latter subthreshold lensed counterparts to confirmed gravitational-wave detections by effectively reducing the noise background while keeping the targeted foreground constant. This chapter explains in details one of the methods being used, namely the GstLAL-based TargetEd Subthreshold Lensing seArch (TESLA) pipeline. We provide an assessment to its performance in searching for potential subthreshold lensed counterparts to superthreshold gravitational waves.

This chapter is structured as follows ¹: In Sec. 7.2, we pose the problem of searching for potential subthreshold lensed counterparts, and introduce the TESLA pipeline’s

¹ The original publication includes an additional section on the overview of how matched-filtering search pipelines work to search for possible gravitational-wave candidates, using GstLAL as an example. It has been omitted here, and readers are asked to refer to Chapter 4.

working principle. In Sec. 7.3, we provide details of a mock data challenge performed to assess the performance of the TESLA pipeline, and compare its effectiveness to alternative proposals in solving the problem posed in the previous section. Finally, Sec. 7.4 summarizes the findings and discusses possible future work to improve the search sensitivity of the TESLA pipeline (at the time of writing).

7.2 The TESLA search method for subthreshold lensed gravitational waves

In this section we will introduce the TargetEd subthreshold Lensing seArch (TESLA) pipeline aiming to search for potential subthreshold lensed counterparts to confirmed superthreshold gravitational waves.

The need for a reduced targeted template bank

As discussed in Chapter 4, a large template bank is used for a general search for gravitational waves to cover a wide parameter space, solely because we have no prior information about the parameters of the gravitational waves we are searching for. However, higher number of templates results in higher trials factors and larger noise background. This will lower the ranking statistics of gravitational wave signal, particularly those being weaker, and caused them to remain un-identified. Hence, we have to develop a way to reduce the nuisance noise background while keeping the targeted foreground constant by reducing the search parameter space, keeping only a subset of templates from the original full template bank.

Deciding which region of the parameter space should be targeted:

Signal sub-space and noise fluctuations

The task upfront now becomes deciding the parameter space that we should be searching in to find subthreshold lensed counterparts to a given targeted superthreshold event. We argue here that there are two major contributing factors: (1) information about the signal sub-space gained from the target superthreshold event, and (2) noise fluctuations in the data, which can lead to finding a candidate event with a template whose parameters differ from those of the target event.

Information about the signal sub-space

Recall from Chapter 6 that strongly-lensed gravitational waves from the same source should have identical waveforms apart from (1) a relative arrival time delay, (2) an overall amplitude scaling factor, and (3) an additional Morse phase factor. That said, the potential subthreshold lensed counterparts we are searching for should

have the same intrinsic parameters (e.g. component masses and spins) as the target superthreshold event. In principle, if we know precisely the underlying parameters for the target event, a single template with the exact same parameters would be ideal to search for its potential subthreshold lensed counterparts. However, the parameters of a gravitational wave are not exactly known, but instead given by best estimates from the posterior probability distribution obtained by Bayesian parameter estimation. Therefore, a good starting point would be to keep templates within the parameter space enclosed by the 90% credible region of the posterior probability distribution of the target event's parameters.

Nevertheless, we argue that the posterior space is insufficient to cover all potential subthreshold lensed counterparts. Bayesian parameter estimation for gravitational waves typically assumes that noise in the data is Gaussian and stationary, which is not true in reality. That said, the posterior probability distribution obtained is for one noise realization only, i.e. the width of the posterior space does not account for noise fluctuations in actual data. Should the superthreshold signal be found at a different time in the data, the resulting posterior probability distribution obtained from Bayesian parameter estimation can be significantly different from the initial one. This argument will be demonstrated in later sections of this chapter.

Noise fluctuations in the data

Should noise in actual data be stationary and Gaussian, the posterior space of the target event would be sufficient to serve as a search sub-space to look for potential subthreshold lensed counterparts. However, noise fluctuations in actual data add complexity, since they can result in the subthreshold signal being found with a template that falls outside this parameter sub-space. Hence, we will need to also consider the effects of noise fluctuations in actual data when deciding which region of the parameter space should be targeted. We do this by injecting subthreshold signals into noisy detector data and identify all the templates that can recover them; this is described in some details below.

An injection campaign accounting for both factors

Strongly-lensed gravitational waves from the same source should have exactly the same waveform, differing only by an overall scaling factor. Hence we can use the posterior samples obtained by Bayesian parameter estimation of the target event to generate possible simulated lensed injections that have similar parameters (i.e.

component masses and spins) and sky location as the target event. To mimic the effect of lensing de-magnification, we reduce the amplitude and hence the SNRs of the injections. This can be done by increasing the source effective distance, as the optimal SNR scales inversely with the source’s effective distance. In detail, we take the posterior samples of the target event and rank them in decreasing order of likelihood. Within a given injection period, we generate, for each posterior sample, one injection with the original optimal SNR, and nine ² additional weaker injections with smaller optimal SNRs by increasing their effective distances, requiring that their SNRs in each detector have to be ≥ 4 to ensure they can be registered as a trigger during the matched-filtering process in the search. These simulated lensed injections represent possible subthreshold lensed counterparts to the superthreshold target event. We then inject these simulated signals into actual data, and use GstLAL to recover ³ them with a general template bank. Because of noise fluctuations, some injections will be found by templates that have parameters significantly different from those within the posterior space of the target event. In the end, we keep only templates that can find these injections, and use them to construct a reduced targeted template bank to search for possible subthreshold lensed counterparts to the target event. Performing the injection campaign allows us to approximate a near-to-optimal targeted template bank taking into account both the information of the signal subspace we gained from the target event (by using the posterior samples to generate simulated lensed injections) and noise fluctuations in actual data. This ensures templates in the reduced bank can identify any potential subthreshold lensed signals while effectively reducing the noise background.

A targeted search to dig up possible lensed candidates

Once a targeted bank is constructed, we again use GstLAL to search through all possible data with the targeted bank to look for potential subthreshold lensed counterparts to the target superthreshold event. As explained in Chapter 4, GstLAL outputs a list of candidate events ranked by their assigned ranking statistics, including FARs and $\ln \mathcal{L}$. It is important to remind readers that the assigned FARs to the candidate events here are not measures of how likely they are to be lensed counterparts to the target event, but rather, as in the full search, we use the FAR to distinguish noise events (false alarms) from real astrophysical signals, whether or

² The number nine is arbitrarily chosen. In principle, one can generate as many injections as one wishes.

³ An injection is considered “recovered” if the corresponding trigger has a False-Alarm-Rate (FAR) $< 1/30$ days.

not they are lensed counterparts to a target event. In this case, we should use the ranking statistics assigned as a priority ranking for follow-up analysis to decide how likely each candidate event is a lensed counterpart to the target event. The details for the follow-up analysis are discussed in [233, 235] and are out of this paper's scope. Readers are reminded that the sole purpose of the TESLA search pipeline is to reduce the nuisance noise background effectively, and in turn up ranking possible subthreshold lensed counterparts to a target superthreshold event, assuming it is strongly lensed. It does not serve the purpose of estimating how likely the found potential candidates are indeed lensed counterparts to the target event.

Figure 7.2.1 summarizes the major steps in the TESLA search pipeline discussed in this section.

7.3 Simulation campaign

We perform a simulation campaign to test the effectiveness of the proposed TESLA pipeline to search for potential subthreshold lensed counterparts to a target superthreshold gravitational wave, assuming it is strongly lensed. Figure 7.3.1 outlines the flow of the simulation campaign. We first prepare a mock data stream with a set of injected lensed signals, one being superthreshold and the other being subthreshold. A general search is then performed using GstLAL with the usual large general template bank ⁴. The general search is expected to recover the superthreshold signal. Bayesian parameter estimation is then performed for the found superthreshold signal, which outputs a set of posterior samples. Then, we apply the TESLA search pipeline to perform an injection campaign and construct a targeted bank to search for the potential subthreshold lensed counterparts to the target event. Finally, we perform another search with GstLAL ⁵ using the targeted bank to see if we can uprank the remaining subthreshold lensed signal that is injected. It has been suggested that extreme template banks ⁶, including

- (1) a single template bank with the template parameters being those of the posterior sample for the target event with the maximum posterior probability,

⁴ The general template bank is composed of several sub-banks targeted different systems. The minimal match of the sub-banks are in general $\geq 97\%$, with certain banks having minimal match $\geq 99\%$, see Table II in [21].

⁵ A re-filtering is required only because (a) PE posteriors correspond to templates that are not in the full template bank, and (b) the results of the search with the full template bank discarded most subthreshold triggers, requiring us to re-run the search pipeline.

⁶ Here, I would like to thank the various reviewers of the short author publications, whom, through the years of review process, help us to strengthen our publication before it was finally published.

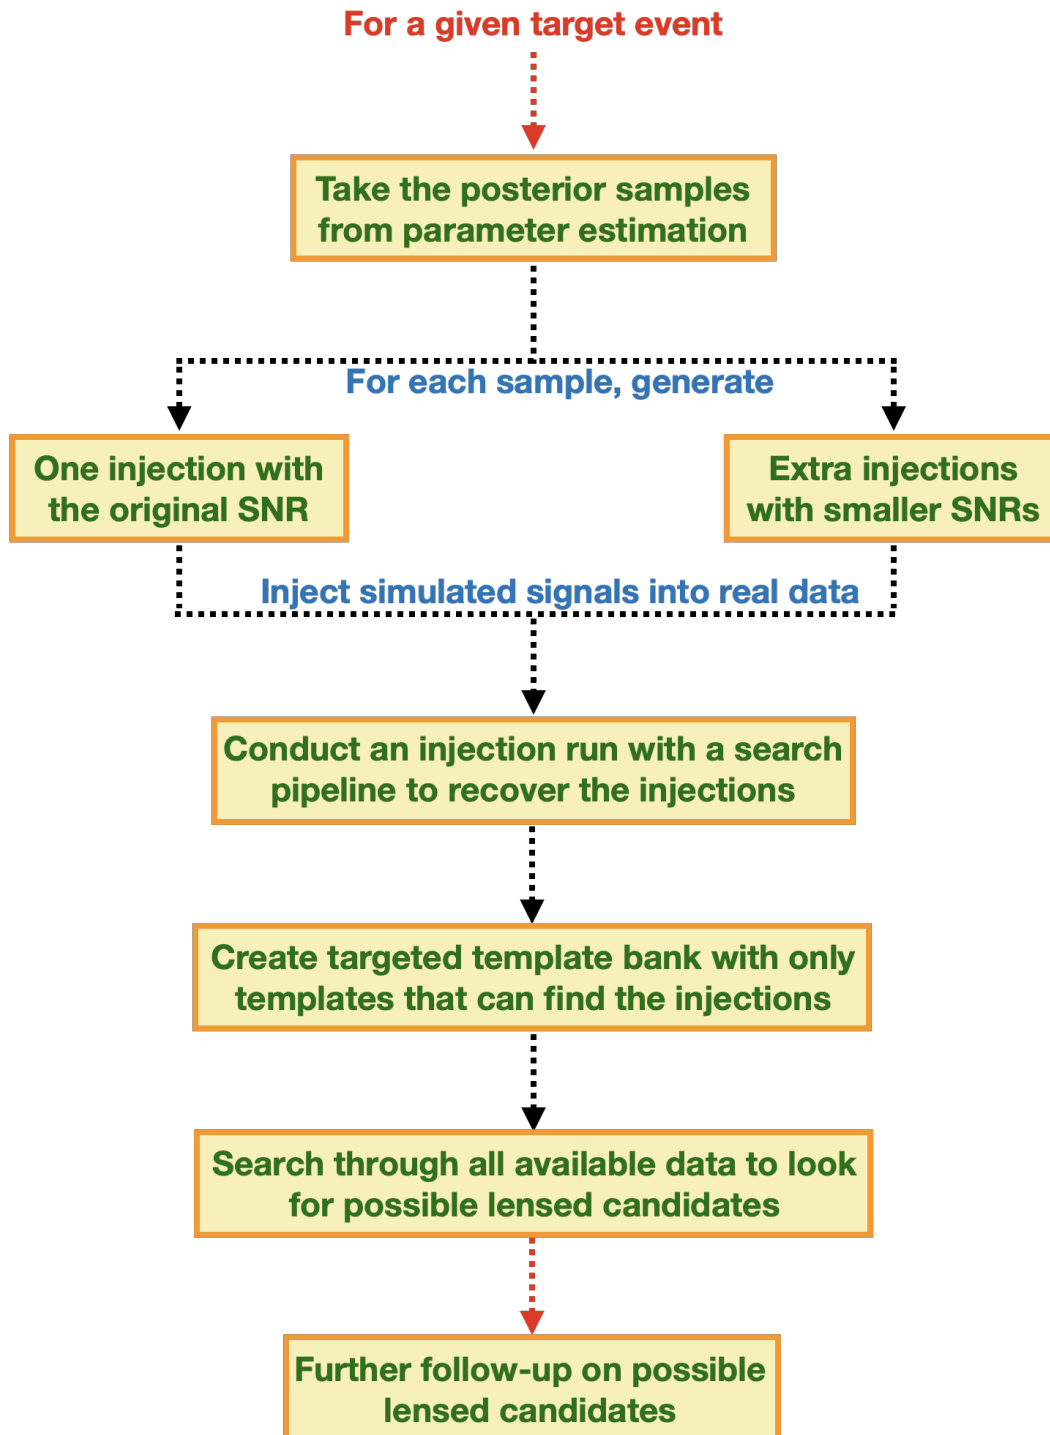


Figure 7.2.1: Workflow of the targeted subthreshold search method (TESLA).

- (2) a PE template bank constructed by only keeping templates that lie within the 90% credible region of the posterior probability distribution for the target event, and

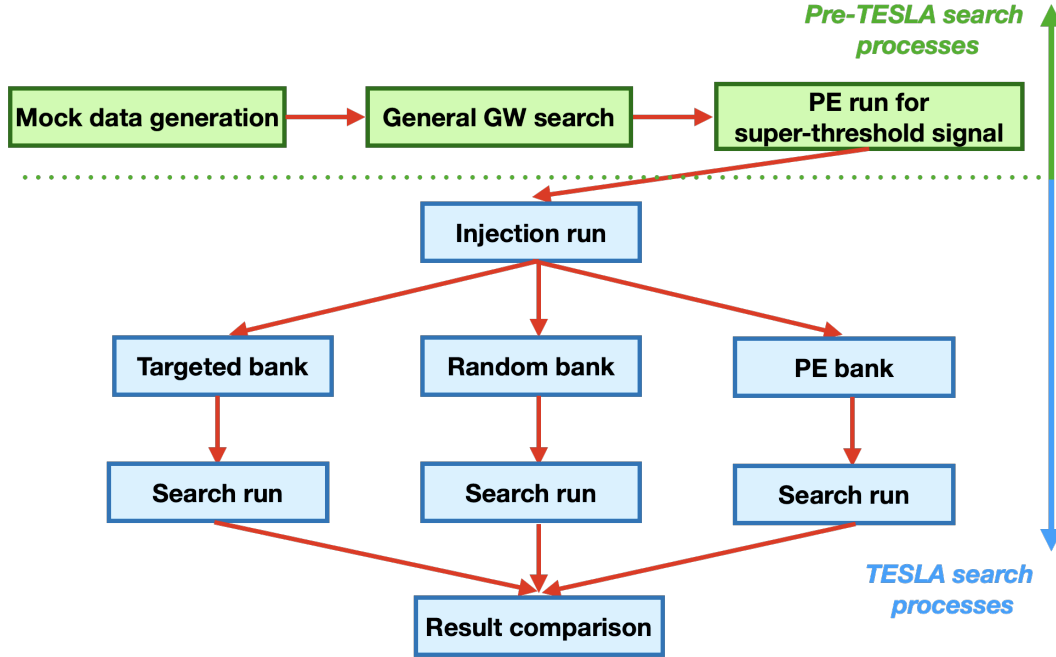


Figure 7.3.1: The analysis flow of the simulation campaign.

- (3) a random template bank constructed by randomly selecting templates from the full bank,

will have higher efficiencies than the targeted bank generated with the proposed TESLA pipeline. We therefore perform extra searches with GstLAL using the proposed banks and compare their performance.

Mock data generation and information

Figure 7.3.2 summarizes the basic information of the mock data stream used for this simulation campaign. For LIGO Hanford, LIGO Livingston and Virgo detector, we generate a ~ 28 -hour-long data stream with Gaussian noise recolored with O3a characteristic power spectral densities (PSDs). We assume no detector downtime⁷, and no times are vetoed. A pair of strongly-lensed gravitational waves simulated following [364] is generated using the SEOBNRv4pseudoFourPN [69] waveform approximant, and is injected into the mock data. The superthreshold signal and subthreshold signal are injected at times shown in figure 7.3.2. Details about the source parameters of the gravitational-wave signal pair are listed in table 7.3.1. We then apply Bilby [49, 305], a Bayesian inference library for gravitational-wave astronomy, to perform parameter estimation (PE) for the superthreshold signal,

⁷ A detector is considered “down” if it is not in observing mode.




Figure 7.3.2: Information about the mock data used for the simulation campaign.

Properties	Superthreshold signal	Subthreshold signal
GPS time	1325932493	1326029051
Magnification μ_i	1.503	-0.595
Distance (Mpc)	2842.00	4518.21
Image type	I	II
Primary mass m_1^{det}	42.0 M_\odot	
Secondary mass m_2^{det}	39.9 M_\odot	
Dimensionless spins	$\chi_{1/2,x} = \chi_{1/2,y} = 0$, $\chi_{1,z} = 0.488, \chi_{2,z} = -0.245$	
Right ascension α	2.939	
Declination δ	0.143	
Inclination ι	2.720	
Polarization Ψ	4.093	
Source redshift z_{source}	0.579	
Lens redshift z_{lens}	0.245	

Table 7.3.1: Information of the injected lensed gravitational-wave pair for the simulation campaign. All properties reported here are measured in the detector frame.

which outputs a set of posterior samples required for applying the TESLA search pipeline.

Applying the TESLA method

Next, we apply the TESLA search pipeline to perform an injection campaign to construct a reduced targeted template bank to search for the remaining subthreshold

lensed counterpart to the superthreshold target event. We generated 5868 simulated lensed injections using the posterior samples obtained from the PE of the superthreshold lensed event. We inject these simulated signals into the mock data and perform another search using GstLAL with the general template bank and try recovering them. 552 templates are rung up by the recovered injections ⁸, and they are used to construct the targeted template bank (see figure 7.3.3). As we can see, even when the noise is almost stationary and Gaussian, subthreshold lensed signals can still be found by templates with parameters very different from those within the posterior space of the superthreshold target event. This demonstrates our earlier argument that the posterior space of the target event itself is insufficient to cover all possible subthreshold lensed counterparts. Finally, we perform another search using GstLAL with the targeted template bank over the same period of mock data to try recovering the remaining injected subthreshold lensed signal. Note that we also included the lensed injection set that was used to determine the templates that we are keeping in the reduced template bank in the search for performance analysis in later sections of the Chapter (see section 7.3). Table 7.3.2 summarizes the search results for the two injected signals using the TESLA targeted template bank. We can

Search results	Superthreshold signal	Subthreshold signal
GPS time	1325932493	1326029051
Rank	1	3
FAR (Hz)	5.37×10^{-21}	4.27×10^{-5}
$\ln \mathcal{L}$	48.63	12.13
Network SNR ρ_{network}	12.20	7.60

Table 7.3.2: Results of the targeted search of the simulation campaign for the two injected lensed signals.

see that (1) The FAR of the subthreshold signal has been reduced by two orders of magnitude, with the log likelihood ratio $\ln \mathcal{L}$ and network SNR ρ_{network} increased. That is, the ranking statistics of the subthreshold signal have been improved. (2) The ranking of the subthreshold signal improves significantly from its previous position of > 100 to the current 3. This means the TESLA search pipeline has successfully upranked the subthreshold signal, and hence made it easier to be identified as a possible gravitational wave for further analysis. We admit that the new FAR of the subthreshold signal still does not pass the usual FAR threshold of 1 in 30 days. This

⁸ Note that this does not mean only 552 injections are recovered. The same template can be used to recover multiple injections. See Table 7.3.6 for the actual number of injections recovered.




Figure 7.3.3: The templates in the original and targeted bank plotted in dark blue and orange respectively on the m_1 - m_2 plane. The best-match template for MGW220111a is indicated by a red star, and the true parameters of MGW220111a is represented by a green diamond. As we can see, even when the noise in the mock data is almost stationary and Gaussian, subthreshold lensed signals can still be found by templates with parameters very different from those within the posterior space of the superthreshold target event. This demonstrates that the posterior space of the target event itself is insufficient to cover all possible subthreshold lensed counterparts.

is primarily due to the observing time being too short. However, we note that the FARs assigned to each candidate here should only be treated as priority ranking for follow-up analysis to determine whether or not (1) they are gravitational waves, and (2) they are lensed counterparts of the target event. The increase in ranking of the subthreshold signal from > 100 to 3 demonstrates that the TESLA search pipeline is effective in reducing unwanted noise background while conserving the desired

foreground, fulfilling its task to uprank potential subthreshold lensed counterparts for a targeted superthreshold event.

Performance comparison with other suggested alternatives

Suggestions have been made that (1) a single template bank, (2) a PE template bank or (3) a random template bank will be more efficient than the targeted template bank constructed with the proposed TESLA pipeline. Here we conduct additional searches using the proposed alternative banks to compare their performance. A random template bank is generated by randomly selecting the same number of templates (i.e. 552 templates) as the targeted template bank. A PE template bank with 81 templates is generated by keeping only templates that lie within the 90% credible region of the posterior probability distribution obtained by Bayesian parameter estimation for the target event. A single template bank in principle should only contain one template with parameters identical to those of the posterior sample with maximum posterior of the target superthreshold event ⁹. Figures 7.3.4, 7.3.5 and 7.3.6 show the distribution of templates in the “single-template” bank, PE bank and random bank respectively.

Recovering the subthreshold lensed signal

Three additional searches using GstLAL are performed over the same period of mock data as the injection run with the random template bank, the PE template bank and “single-template” bank respectively, in order to recover the subthreshold injected signal. Tables 7.3.3, 7.3.4 and 7.3.5 summarize the search results for the two injected signals.

Search results	Superthreshold signal	Subthreshold signal
GPS time	1325932493	1326029051
Rank	1	> 100
FAR (Hz)	3.07×10^{-14}	1.54×10^{-2}
$\ln \mathcal{L}$	28.0	-2.14
Network SNR ρ_{network}	10.45	7.21

Table 7.3.3: Results of the search for the two injected lensed signals using the random template bank.

⁹ For practical reasons, we use a bank with 100 templates having component masses within $\pm 0.1 M_{\odot}$ from those of the posterior sample with maximum posterior to mimic the single template bank.




Figure 7.3.4: The templates in the original and “single-template” bank plotted in dark blue and orange respectively on the m_1 - m_2 plane. The “single-template” bank is a bank with 100 templates having component masses within $\pm 0.1 M_\odot$ from those of the posterior sample with maximum posterior to mimic the single template bank.

Search results	Superthreshold signal	Subthreshold signal
GPS time	1325932493	1326029051
Rank	1	7
FAR (Hz)	3.61×10^{-5}	9.05×10^{-5}
$\ln \mathcal{L}$	48.13	12.69
Network SNR ρ_{network}	12.53	7.663

Table 7.3.4: Results of the search for the two injected lensed signals using the PE template bank.




Figure 7.3.5: The templates in the original and PE bank, plotted in dark blue and orange respectively on the m_1 - m_2 plane. The best-match template for GW220111a is indicated by a red star, and the true parameters of GW220111a are represented by a green diamond. The purple curve represents the boundary to the 90% credible region of the posterior probability distribution for GW220111a. The PE bank is generated by keeping only templates that lie within the 90% credible region of the posterior probability distribution, containing only 81 templates.

From the results, we see that (1) the “single-template” bank fails to even register the subthreshold signal as a trigger in the first place during the matched-filtering process, (2) the PE template bank successfully upranks the subthreshold signal to a rank 7 candidate and improves its ranking statistics, but its performance is not as good as compared to that using the TESLA search pipeline, and (3) the random template bank fails to improve the ranking and the ranking statistics of the subthreshold event. This means that the targeted foreground is affected by the reduction in the number




Figure 7.3.6: The templates in the original and random bank plotted in dark blue and orange respectively on the m_1 - m_2 plane. The best-match template for GW220111a is indicated by a red star, and the true parameters of GW220111a are represented by a green diamond. The random bank contains the same number of templates, i.e. 552 templates, as the targeted template bank, and they are randomly selected from the original template bank.

of templates for the random template bank. It is therefore evident that the random bank is not suitable to search for potential subthreshold lensed gravitational waves.

Search results	Superthreshold signal	Subthreshold signal
GPS time	1325932493	...
Rank	1	...
FAR (Hz)	6.11×10^{-6}	...
$\ln \mathcal{L}$	24.11	...
Network SNR ρ_{network}	12.49	...

Table 7.3.5: Results of the search for the two injected lensed signals using the “single-template” bank.

Simulated lensed injections recovery

To further compare the performance of the banks, we analyse the change in number of lensed injections recovered ¹⁰ using the four proposed banks as compared to using the general template bank. Table 7.3.6 summarizes the findings.

Injections	General	TESLA	Random	PE	Single
Total	5868	5868	5868	5868	5868
Found	1793	1959	402	299	1076
Missed	4075	3909	5466	5569	4792
Found % change	-	+9.26%	-77.5%	-80.3%	-40.0%

Table 7.3.6: Number of injections found and missed during the search of mock data using the general template bank, TESLA targeted template bank, random template bank, PE template bank and “single-template” bank respectively.

The targeted template bank constructed using the TESLA search pipeline succeeds in recovering more lensed subthreshold injections than the other banks. The “single-template” bank, the PE template bank and the random template bank miss more lensed subthreshold injections. The random template bank is expected not to give a satisfactory performance in recovering the lensed subthreshold injections. While the lensed subthreshold injections are generated using the exact same parameters as the posterior samples of the target event (i.e. the injections should all have similar parameters as to the templates in the PE template bank), the PE template bank misses even more injections as compared to using the general template bank ¹¹. This again demonstrates our argument that constructing the targeted template bank

¹⁰ As before, an injection is considered “recovered” if the corresponding trigger has a FAR < 1/30 days.

¹¹ In fact, it misses even more injections than the random bank, but this should not be alarming. Given that the injections are subthreshold, they are more likely to be recovered by templates with very different parameters than their true parameters. The random bank, while being completely

solely by considering the posterior signal sub-space of the superthreshold target event is insufficient. In this simulation campaign, we are simply fortunate that the PE bank can recover the injected subthreshold lensed signal. Should the injected subthreshold signal be even weaker, or should it be injected at a time at which noise is very different from that around the superthreshold signal, the PE template bank is more likely to miss it. On the other hand, the targeted template bank created using the TESLA pipeline is more likely to recover it since the bank is constructed by considering both information about the signal sub-space gained from the target event as well as noise fluctuations in the data.

Sensitive range at different FAR threshold

Finally, we use the 5868 lensed subthreshold injections¹² to evaluate the sensitive range¹³ at different combined FAR threshold for each template bank. Figure 7.3.7 shows the percentage changes in sensitive range v.s. FAR curves obtained using the alternative banks as compared to that using the full template bank for lensed subthreshold signals that are similar to the target superthreshold event. We can see that the percentage-change curve representing results using the targeted bank constructed by the TESLA search pipeline is above that of the full template bank, showing improvement in terms of sensitivity towards MGW220111a-alike (lensed) subthreshold signals. Meanwhile, the same curves for the random bank, the PE bank and the “single-template bank” are below that of the full template bank, showing that the sensitivity towards MGW220111a-alike (lensed) signals is worsened. This further demonstrates that the targeted template bank created using the TESLA search pipeline has the best performance among the four banks to search for potential subthreshold lensed gravitational waves for a target superthreshold event.

Summary of results

To sum up the results presented above, in this simulation campaign we investigated four proposed banks to search for possible subthreshold lensed counterparts of a given superthreshold gravitational wave. Three of the cases are found to be inferior

random, covers a much wider parameter space than the PE bank, and hence have a higher chance in recovering the subthreshold injections.

¹² These are the same injections used in the simulation campaign to create the reduced template bank.

¹³ The sensitive range is the distance out to which we may identify gravitational waves averaged over relevant parameters including sky location and binary orientation. Note that in this analysis we assumed the injections are not lensed, i.e. they have magnification $\mu = 1$.




Figure 7.3.7: (Top panel) The sensitive range v.s. FAR for MGW220111a-alike signals using the full template bank (black), targeted template bank (orange), “single-template” bank (blue), PE template bank (green) and random template bank (red) respectively. The shaded band for each curves represent the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. The percentage-change curve (orange) representing results using the targeted bank constructed by the TESLA search pipeline is above that of the full template bank, showing improvement in terms of sensitivity towards MGW220111a-alike (lensed) subthreshold signals. The vertical striped line shows the typical FAR threshold for triggers below which we consider as possible lensed candidates.

in performance compared with the TESLA bank, namely the single-template bank, the PE template bank and the random template bank. The results show that none of the three alternative cases can outperform an intermediate template bank created based on the TESLA pipeline in terms of search sensitivity and effectiveness.

It should be noted that in this simulation campaign, we considered the case where lensing creates a pair of repeated gravitational-wave signals from the same source, separated by roughly a day (1.11 days). In practice, the relative time delay between repeated signals can range from minutes to months for galaxy lenses. We will need to perform the injection campaign over a longer time range with a larger number of injections. This will result in an increase in size for the reduced template bank, and may affect the performance of the reduced template bank. The ranking statistics,

in particular the FAR of the triggers, will also be affected based on the number of templates we have in the reduced template bank. Chapter 11 discuss some of the possible work to fine tune the selection procedure for templates included in the targeted template bank in order to find the optimal balance between coverage and sensitivity.

7.4 Concluding Remarks

At the time of writing, the LVK collaboration has just published the first full-scale analysis to search for lensing signatures of gravitational waves within the first half of LIGO/Virgo third observing run O3a [23], and concluded that there is not yet any compelling evidence for gravitational lensing of gravitational waves. One featured analysis in the paper explores the possibility of strong lensing producing magnified superthreshold gravitational-wave signals, and de-magnified subthreshold copies that have insufficient significance and remain un-identified as detections. Two independent search methods were applied to search for the latter potential subthreshold lensed signals, one being the GstLAL-based TargetEd subthreshold Lensing seArch (TESLA) pipeline.

In a general search for gravitational waves, a large template bank covering a wide parameter space is used since we have no prior information regarding the parameters of gravitational waves we are searching for. The large number of templates used contributes a high trials factors. This may bury potential subthreshold (lensed) gravitational waves in the huge noise background. To search for possible subthreshold lensed counterparts to superthreshold confirmed gravitational waves, we need to reduce the noise background while keeping the targeted foreground constant. In other words, we want to lower the noise background by tactically focusing a particular region in the parameter subspace, while keeping the targeted foreground constant, and hence upranking any potential subthreshold lensed candidates to the superthreshold target events.

In this Chapter we explain the methodology of the TESLA pipeline in detail, and demonstrate that the TESLA pipeline can efficiently search for possible subthreshold lensed counterparts to confirmed superthreshold gravitational-wave detections.

The TESLA pipeline fulfils the task by conducting an injection campaign. It prepares simulated lensed injections based on posterior samples obtained from Bayesian parameter estimation of the superthreshold target event, such that they have similar intrinsic parameters as the target event, but with varying effective distances and

hence weaker amplitudes to mimic the de-magnifying effect caused by gravitational lensing. These injections are then injected into actual data ¹⁴ and a GstLAL search is performed using the general template bank to recover these injections. Templates that can find the injections are used to construct a targeted template bank, which is then used to perform another GstLAL search to look for possible subthreshold lensed counterparts to the target event, should it be strongly lensed. We argue that the TESLA search pipeline can generate a template bank that performs better than alternatives to search for these potential subthreshold signals, as it accounts for both information about the signal-subspace gained from the target event, as well as noise fluctuations in actual data.

To assess the performance of the TESLA search pipeline, we conducted a simulation campaign in which we simulated LHO, LLO, and Virgo data streams with Gaussian noise recolored with O3a representative power spectral densities (PSDs) and a pair of lensed signals, one being superthreshold and the other being subthreshold. We first perform a GstLAL search using the general template bank to recover the superthreshold signal, and perform Bayesian parameter estimation to generate a set of posterior samples. Then, we use the TESLA pipeline and try to recover the remaining subthreshold lensed signal from the mock data. Our results show that the TESLA pipeline can effectively uprank the subthreshold signal, improving the probability that it will be identified as a gravitational wave, and with further analysis, a lensed counterpart to the target superthreshold event.

We also compare the performance of the targeted template bank constructed with the TESLA search pipeline to suggested alternative template banks: (1) a single template bank with the template parameters being those of the posterior sample for the target event with the maximum posterior probability, (2) a PE template bank constructed by keeping only templates from the general bank that lie within the 90% credible region of the posterior space for the target event, and (3) a random template bank constructed by randomly selecting templates from the general template bank. We show, by considering their performance in (1) recovering the injected subthreshold lensed signal, (2) recovering the simulated lensed injections and (3) their sensitive range for gravitational waves that are similar to the target events, that the targeted template bank constructed using the TESLA search pipeline outperforms the other three alternative banks. In fact, the results show that one would not expect additional improvement when further narrowing the template bank.

¹⁴ Note that the results in this paper make use of Gaussian simulated data, but for the actual subthreshold search, these injections are made into actual data.

The search sensitivity of the TESLA search pipeline can be further improved. For instance, since we are looking for lensed counterparts of targeted events, using the target's sky location, we should be able to set a consistent range for the difference in arrival time and phase between participating detectors for the lensed counterparts. This will be discussed in a future paper under development. Also, the selection procedure for templates included in the targeted template bank may require further tuning to find the optimal balance between coverage and sensitivity.

This method has been applied to subsequent searches for subthreshold lensed events, in LVK papers [23, 32].

Chapter 8

SEARCH FOR LENSING SIGNATURES IN THE FIRST HALF OF LVK’S THIRD OBSERVING RUN

Note: This Chapter is an adaptation of the LIGO-Virgo-KAGRA collaboration-wide publication:

LIGO Scientific and Virgo Collaborations, R. Abbott, ...(723 authors)..., **Alvin K. Y. Li**, ...(655 authors)..., “Search for Lensing Signatures in the Gravitational-Wave Observations from the First Half of LIGO–Virgo’s Third Observing Run”, *Astrophys. J.* 923.1 (2021), doi: 10.3847/1538-4357/ac23db

Alvin K. Y. Li led the analysis for the search for sub-threshold lensed gravitational waves and edited the corresponding section in the LIGO-Virgo-KAGRA collaboration-wide paper (Section 5.3). The other related sections in this chapter are contributions from co-authors.

8.1 Introduction

In this chapter we perform a comprehensive lensing analysis of data from the first half of the third LIGO–Virgo observing run, called O3a for short, focusing on compact binary coalescence (CBC) signals. We begin by outlining the expected rate of strongly lensed events. Strong lensing is rare, but magnified signals enable us to probe a larger comoving volume, thus potentially giving us access to more sources [120, 224, 274, 281, 302, 308, 323–325]. We forecast the lensed event rates using standard lens and black hole population models (Sec. 8.3). These expected rates are subject to some astrophysical uncertainty but are vital to interpreting our search results in later sections.

The rate of lensing can also be inferred from the stochastic GW background (SGWB) [81, 82, 266]. Thus, we use the non-observation of strong lensing and the stochastic background to constrain the binary black hole (BBH) merger-rate density and the rate of lensing at high redshifts.

Since lensing magnification biases the redshift z estimation, the inferred gravitational wave (GW) luminosity distance and source mass measurements (source

frame masses differ from detector frame masses by a factor of $(1+z)^{-1}$ can also be biased. This could lead to observations of apparently high-mass (or low-mass, when de-magnified) binaries [74, 75, 120, 186, 281]. Therefore, we analyze several LIGO–Virgo detections with unusually high masses under the alternative interpretation that they are lensed signals from lower-mass sources which have been magnified (Sec. 8.4).

We then move on to search for signatures of lensing-induced multiple images, which should appear as repeated similar signals, magnified and with waveform differences determined by the image type [119, 154], separated in time by minutes to months (or even years). Consequently, if an event pair is strongly lensed, we expect to infer consistent parameters for both events [186, 188].

We search for these multiple images by first comparing the posterior overlap between pairs of events occurring during the O3a period as reported in [21] (Sec. 8.5). After identifying a list of candidates from the posterior-overlap analysis, we follow these up with more computationally expensive but more accurate joint-parameter estimation (PE) procedures (Sec. 8.5). Next, we perform a targeted search for previously undetected counterpart images of known events in Sec. 8.5, images that could have fallen below the threshold of previous wide-parameter space CBC searches [as discussed in 122, 223, 247]. Finally, we search for microlensing induced by point-mass lenses in the intermediate and low mass range, including wave-optics effects (Sec. 8.6).

Several searches for GW lensing signatures have already been performed in the data from the first two observing runs O1 and O2 [122, 186, 223, 233, 247, 288], including strong lensing and microlensing effects. We will discuss these previous studies in the appropriate sections. Given the growing interest in GW lensing and the existing forecasts, an analysis of the most recent GW observations for lensing effects is now timely.

Results of all analyses in this paper and associated data products can be found in LIGO Scientific Collaboration and Virgo Collaboration [228]. Gravitational-wave strain data and posterior samples for all events from GWTC-2 are available [181] from the Gravitational Wave Open Science Center [34].

8.2 Data and events considered

The analyses presented here use data taken during O3a by the Advanced LIGO [1] and Advanced Virgo [36] detectors. O3a extended from 2019 April 1 to 2019

October 1. Various instrumental upgrades have led to more sensitive data, with median binary neutron star (BNS) inspiral ranges [43] increased by a factor of 1.64 in LIGO Hanford, 1.53 in LIGO Livingston, and 1.73 in Virgo compared to O2 [21]. The duty factor for at least one detector being online was 97%; for any two detectors being online at the same time it was 82%; and for all three detectors together it was 45%. Further details regarding instrument performance and data quality for O3a are available in Abbott et al. [21], Buikema et al. [79].

The LIGO and Virgo detectors used a photon recoil based calibration [83, 207, 357] resulting in a complex-valued, frequency-dependent detector response with typical errors in magnitude of 7% and 4 degrees in phase [37, 335] in the calibrated O3a strain data.

Transient noise sources, referred to as glitches, contaminate the data and can affect the confidence of candidate detections. Times affected by glitches are identified so that searches for GW events can exclude (veto) these periods of poor data quality [3, 15, 131, 163, 275]. In addition, several known noise sources are subtracted from the data using information from witness auxiliary sensors [128, 145].

Candidate events, including those reported in Abbott et al. [21] and the new candidates found by the searches for sub-threshold counterpart images in Sec. 8.5 of this paper, have undergone a validation process to evaluate if instrumental artifacts could affect the analysis; this process is described in detail in Sec. 5.5 of [131]. This process can also identify data quality issues that need further mitigation for individual events, such as subtraction of glitches [114] and non-stationary noise couplings [352], before executing PE algorithms. See Table V of Abbott et al. [21] for the list of events requiring such mitigation.

The GWTC-2 catalog [21] contains 39 events from O3a (in addition to the 11 previous events from O1 and O2) with a false-alarm rate (FAR) below two per year, with an expected rate of false alarms from detector noise less than 10% [21]. We neglect the potential contamination in this analysis. These events were identified by three search pipelines: one minimally modeled transient search cWB [212–214, 216, 217] and the two matched-filter searches GstLAL [183, 250, 309] and PyCBC [42, 43, 123, 276, 351]. Their parameters were estimated through Bayesian inference using the LALINFERENCE [356] and Bilby [49, 305, 326] packages. Both the matched-filter searches and PE use a variety of CBC waveform models which generally combine knowledge from post-Newtonian theory, the effective-one-body formalism and numerical relativity [for general introductions to these approaches,

see 61, 125, 287, 314, and references therein]. The analyses in this paper rely on the same methods, and the specific waveform models and analysis packages used are described in each section.

Most of the 39 events from O3a are most probably BBHs, while three (GW190425, GW190426_152155, and GW190814) have component masses below $3M_{\odot}$ [12, 19, 21], thus potentially containing a neutron star. We consider these 39 events in most of the analyses in this paper, except in the magnification analysis (Sec. 8.4), which concerns only six of the more unusual events, and the microlensing analysis (Sec. 8.6), which focuses on the 36 clear BBH events only.

Specifically, we use the following input data sets for each analysis. The magnification analysis in Sec. 8.4 and posterior-overlap analysis in Sec. 8.5 start from the Bayesian inference posterior samples released with GWTC-2 [181]. The joint-PE analyses in Sec. 8.5 and microlensing analysis in Sec. 8.6 reanalyze the strain data in short segments around the event times, available from the same data release, with data selection and noise mitigation choices matching those of the PE analyses in Abbott et al. [21]. In addition, the searches for sub-threshold counterpart images in Sec. 8.5 cover the whole O3a strain data set, using the same data quality veto choices as in Abbott et al. [21] but a strain data set consistent with the PE analyses: the final calibration version of LIGO data [335] with additional noise subtraction [352].

8.3 Lensing statistics

In this section, we first forecast the number of detectable strongly lensed events (Sec. 8.3). Then, we infer upper limits on the rate of strongly lensed events using two different methods; the first uses only the non-detection of resolvable strongly lensed BBH events (Sec. 9.5), while the second leverages additionally the non-observation of the SGWB (Sec. 8.3) [20, 84]. Since the background would originate from higher redshifts, this second method complements the first method.

Throughout this section, we model the mass distribution of BBHs following the results for the POWER LAW + PEAK MODEL of [22]. We consider two distinct models of the BBH merger rate density. Model A brackets most of the population synthesis results [65, 148, 273, 311] corresponding to Population I and II stars while Model B assumes the Madau–Dickinson ansatz [239] where the rate peaks at a particular redshift. For consistency with previous analyses [e.g., 20], we take the Hubble constant from Planck 2015 observations to be $H_0 = 67.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [41]. Detailed discussion on both models and their respective parametrization is given

in Appendix 8.A. The obtained rates are subject to uncertainty because of their dependence on the merger rate density, which is model-dependent and only partially constrained. They are nevertheless vital to interpreting our search results in later sections (see Sec. 8.5).

Strong lensing rate

We predict the rate of lensing using the standard methods outlined in the literature [224, 267, 274, 281, 364, 372], at galaxy and galaxy-cluster lens mass scales. To model the lens population, we need to choose a density profile and a mass function. We adopt the singular isothermal sphere (SIS) density profile for both galaxies and clusters. Moreover, we use the velocity dispersion function (VDF) from the Sloan Digital Sky Survey [SDSS; 102] for galaxies and the halo mass function from Tinker et al. [342] for clusters which have been used in other lensing studies as well [e.g., 282, 302]. The SIS profile can well describe galaxies. However, the mass distribution of clusters tends to be more complicated. Nevertheless, [302] have demonstrated that the SIS model can reproduce the lensing rate predictions from a study of numerically simulated cluster lenses. Thus, we adopt the same model. Under the SIS model, we obtain two images with different magnifications and arrival times. The rate of strong lensing is

$$\begin{aligned} \mathcal{R}_{\text{lens}} = & \int \frac{dN(M_h, z_l)}{dM_h} \frac{dD_c}{dz_l} \frac{\mathcal{R}_m(z_m)}{1+z_m} \frac{dV_c}{dz_m} \sigma(M_h, z_l, z_m, \rho, \rho_c) \\ & \times p(\rho|z_m) d\rho dz_m dz_l dM_h, \end{aligned} \quad (8.1)$$

where $dN(M_h, z_l)/dM_h$ is the differential comoving number density of lensing halos in a halo mass shell at lens redshift z_l , D_c and V_c are the comoving distance and volume, respectively, at a given redshift, $\mathcal{R}_m(z_m)$ is the total comoving merger rate density at redshift z_m , $(1+z_m)$ accounts for the cosmological time dilation, $p(\rho | z_m)$ is the distribution of signal-to-noise ratio (SNR) at a given redshift, and σ is the lensing cross-section (Appendix 8.A). Throughout this section and in Sec. 8.3 we choose a network SNR threshold $\rho_c = 8$ as a point estimator of the detectability of GW signals. We find it to be consistent with the search results in Abbott et al. [21] and in Sec. 8.5, and we estimate its impact to be subdominant with respect to other source of uncertainties.

In Table 8.3.1, we show our estimates of the relative rate of lensing assuming different models (Models A and B) for the merger rate density. The results are shown separately for galaxy-scale (G) and cluster-scale (C) lenses. Furthermore, these rates are calculated for events that are doubly lensed and for two cases: when

Table 8.3.1: Expected fractional rates of observable lensed double events at current LIGO–Virgo sensitivity.

Merger Rate Density Model	Galaxies ($\times 10^{-4}$)		Galaxy Clusters ($\times 10^{-4}$)	
	R_D	R_S	R_D	R_S
A	0.9–4.4	2.9–9.5	0.4–1.8	1.4–4.1
B	1.0–23.5	2.5–45.2	0.7–10.9	1.6–19.9

This table lists the relative rates of lensed double events expected to be observed by LIGO–Virgo at the current sensitivity where both of the lensed events are detected (R_D) and only one of the lensed events is detected (R_S) above the SNR threshold. For Model A, the range corresponds to the bracketing function (see Appendix 8.A) and for Model B, the rates encompass a 90 per cent credible interval. We show the rate of lensing by galaxies ($\sigma_{\text{vd}} = 10\text{--}300 \text{ km s}^{-1}$) and galaxy clusters ($\log_{10}(M_{\text{halo}}/M_{\odot}) \sim 14\text{--}16$) separately. Besides their usage for forecasts, the fraction of lensed events allows us to interpret the prior probability of the strong lensing hypothesis, which we require to identify lensed events confidently.

only a single event (i.e., the brighter one) is detected (S), and when both of the doubly lensed events are detected (D). The expected fractional rate of lensing (lensed to unlensed rate), which will be necessary for the multi-image analyses (Sec. 8.5), ranges from $O(10^{-3}\text{--}10^{-4})$, depending on the merger rate density assumed. We estimate the fractional rate of observed double (single) events for galaxy-scale lenses in the range 0.9–4.4 (2.9–9.5) when using Model A for the merger rate density. Similarly, for cluster-scale lenses, the fractional rate is estimated to be in the range of 0.4–1.8 (1.4–4.1) much rarer than the rates at galaxy scales. These estimates suggest that observing a lensed double image is unlikely at the current sensitivity of the LIGO–Virgo network of detectors. Nevertheless, at design sensitivity and with future upgrades, standard forecasts suggest that the possibility of observing such events might become significant [224, 267, 274, 281, 364, 372]. Our lensing rates are consistent with those predicted for singular isothermal ellipsoid models [e.g., 281, 364, 372]. The main uncertainty in the rate estimates derives from the uncertainties in the merger-rate density at high redshift.

Depending on the specific distribution of lenses and the source population, the time delays between images can change. Models favoring galaxy lensing produce minutes to perhaps months of time delay, while galaxy cluster lensing can produce time delays up to even years. However, the time delay distribution for galaxy cluster lenses is more difficult to model accurately, owing to the more complex lensing

morphology.

Since the merger rate density at high redshift is observationally constrained only by the absence of the SGWB, these rates are subject to uncertainty. Nevertheless, standard theoretical models will still produce useful forecasts. We will later refer to these rate estimates in the relevant sections (see Sec. 8.5).

Implications from the non-observation of strongly lensed events

Motivated by the absence of evidence for strong lensing (Sec. 8.5), we assume that no strong lensing has occurred, in order to constrain the merger rate density at high redshift. We use the standard constraints on the merger rate density at low redshift from the LIGO–Virgo population studies [22]. We assume the Madau–Dickinson form for the merger rate density (Model B). This model’s free parameters include the local merger rate density, the merger rate density peak, and the power-law slope. The non-observation of lensing constrains the merger-rate density at high redshift, which is unconstrained by the low-redshift observations alone (Fig. 8.3.1). These lensing constraints are complementary to the current strictest high-redshift limits obtained through SGWB non-observation [20].

Constraints from stochastic background

We can also constrain the redshift evolution of the merger rate density from the reported non-observation of the SGWB from BBHs [20, 84]. This, in turn, provides constraints on the relative abundance of distant mergers, which are more likely to undergo lensing. Thus, the non-observation of the SGWB can inform the estimate of the probability of observing lensed BBH mergers [81, 266].

Following [81], we forecast constraints on the merger rate density in O3 using up-to-date constraints on the mass distribution and redshift evolution of BBH mergers obtained from the latest detections [10, 11, 21, 22], as well as those inferred from current upper limits on the SGWB, given its non-observation [20].

While the measured parameters for each merger (redshifts, source masses) are potentially biased by lensing, as discussed in Sec. 8.4, we express all quantities as functions of non-biased merger redshift z_m and chirp mass \mathcal{M} [81] for consistency with other sections. However, following Buscicchio et al. [81], we do not assume as prior information that lensing is not taking place. Instead, we include the magnification bias self-consistently in the analysis, by imposing population constraints in apparent masses and redshifts.




Figure 8.3.1: Merger rate density as a function of redshift based on the GWTC-2 results without lensing constraints (blue) and with lensing (red) included in LIGO–Virgo detections. We show the results for galaxy-scale lenses (G) and cluster-scale lenses (C) separately. Furthermore, S (or D) correspond to doubly lensed events where single (or double) events are detected. Because lensed detections occur at higher redshifts than unlensed events, their non-observation can be used to constrain mergers at higher redshifts. The results without lensing do not include constraints derived from the absence of a SGWB.

We model the differential lensing probability following [120]. The differential merger rate in a redshift and magnification shell is

$$\frac{d^2\mathcal{R}}{dz_m d\ln\mu} = \frac{dP(\mu | z_m)}{d\ln\mu} \frac{4\pi D_c^2(z_m)}{H_0(1+z_m)E(z_m)} \times \int dm_1 dm_2 \frac{d^3\mathcal{R}_m(z_m)}{dm_1 dm_2 dz_m} p(\rho > \rho_c | m_1, m_2, z_m, \mu), \quad (8.2)$$

where $d^3\mathcal{R}_m(z_m)/dm_1 dm_2 dz_m$ is the differential merger rate density; $p(\rho > \rho_c | m_1, m_2, z_m, \mu)$ provides the probability of observing mergers with source masses m_1, m_2 , redshift z_m , and magnified by a factor μ above a fixed network SNR threshold $\rho_c = 8$, integrated over the population distribution of source parameters; the factor $4\pi D_c^2(z_m)/[H_0(1+z_m)E(z_m)]$ gives the comoving volume of a redshift shell in an expanding Universe (taking into account the redshifted rate definition with respect to the source frame); and $dP(\mu | z_m)/(d\ln\mu)$ is the lensing probability. However, as




Figure 8.3.2: Cumulative fraction of lensed detectable BBH mergers at any redshift with magnification greater than μ , constrained by the non-observation of the SGWB. The solid line shows the value obtained from the median BBH merger rate density posterior. The shaded region corresponds to the 90% credible interval. Fewer than 1 in 10^3 events are expected to be lensed with magnification $\mu > 2$, on average. Significantly higher magnifications (e.g., $\mu > 30$) are suppressed by a further factor of 10. The results here show the probability of *observing* an event above a given magnification, which includes the merger-rate density and magnification bias information.

noted by Dai et al. [120], the differential magnification probability at $0.9 < \mu < 1.1$ and $z_m < 2$ is affected by relative uncertainties up to 40%. We therefore consider magnified detections only ($\mu > 1$), which are subject to less uncertainty, and normalize our results accordingly. We then integrate the differential merger rate (Eq. 8.2) over redshift and magnifications in $[\mu, \mu_{\max}]$ and divide it by the total rate of magnified detections. By doing so, we obtain the cumulative fraction of detected lensed events at any redshift with magnifications larger than μ .

We show the result in Fig. 8.3.2. We find the observation of lensed events to be unlikely, with the fractional rate at $\mu > 2$ being $4.9^{+1.7}_{-1.3}$. More significantly magnified events are even more suppressed, with a rate of $3.5^{+0.6}_{-0.4}$ at $\mu > 30$. These estimates suggest that most binary mergers that we observe are not strongly lensed. However, as projected in Buscicchio et al. [81], Mukherjee et al. [266], at design sensitivity, the same probability will be enhanced, as a widened horizon will probe the merger

rate density deeper in redshift.

Comparing the above predictions with the expected fractional rates R_S of single lensed detections with Model B in Table. 8.3.1, the predictions agree within a factor of 5 for the relative rate of lensing. The differences are due to a different underlying lens model and partly to the inclusion of de-magnified events in Sec. 8.3.

8.4 Analyzing high-mass events

If a GW signal is strongly lensed, it will receive a magnification μ defined such that the GW amplitude increases by a factor $|\mu|^{1/2}$ relative to an unlensed signal. The luminosity distance inferred from the GW observation will be degenerate with the magnification such that the inferred luminosity distance

$$D_L^{\text{inferred}} = \frac{D_L}{\sqrt{|\mu|}}. \quad (8.3)$$

Because of this degeneracy, lensing biases the inferred redshift and thus the source masses. Consequently, the binary appears to be closer than it truly is, and it appears to be more massive than it truly is.

Broadhurst et al. [74, 75, 77] argued that some of the relatively high-mass LIGO–Virgo events could be strongly lensed GWs from the lower-mass stellar black hole population observed in the electromagnetic bands. However, the expected strong lensing rates and the current constraints on the merger-rate density, based on the absence of a detectable SGWB, disfavor this interpretation [81, 82, 120, 186, 224, 274, 281] compared to the standard interpretation of a genuine unlensed high-mass population [10, 22, 211, 306]. Hence, in the absence of more direct evidence, such as identifying multiple images within LIGO–Virgo data (Sec. 8.5), it is difficult to support the lensing hypothesis purely based on magnification considerations. Nevertheless, it is informative to analyze the degree to which the lensed interpretation would change our understanding of the observed sources.

Under the strong lensing hypothesis \mathcal{H}_{SL} , the GW would originate from a well-known, intrinsically lower-mass population, and the LIGO–Virgo observations have been biased by lensing. Using such a mass prior, we infer the required magnification and corrected redshift and component masses under \mathcal{H}_{SL} . The posterior distribution of the parameters is [288]

$$p(\mu, \vartheta | d, \mathcal{H}_{\text{SL}}) \propto p(d | \vartheta) p(\vartheta | \mu, \mathcal{H}_{\text{SL}}) p(\mu | \mathcal{H}_{\text{SL}}), \quad (8.4)$$

where we distinguish the *apparent* parameters of the waveform received at the detector ϑ , which differ from the intrinsic parameters θ due to bias by lensing

Table 8.4.1: Inferred properties of selected O3a events under the lensing magnification hypothesis.

Event name	$m_1 [M_\odot]$	$m_2 [M_\odot]$	z	μ
GW190425	$1.42^{+0.16}_{-0.12}$	$1.27^{+0.12}_{-0.15}$	$0.3^{+0.1}_{-0.1}$	68^{+163}_{-44}
GW190426_152155	$1.89^{+0.40}_{-0.55}$	$0.90^{+0.25}_{-0.40}$	$1.3^{+0.5}_{-0.2}$	497^{+452}_{-272}

Event name	$m_1^{50} (m_1^{65}) [M_\odot]$	$m_2^{50} (m_2^{65}) [M_\odot]$	$z^{50} (z^{65})$	$\mu^{50} (\mu^{65})$
GW190521	$43^{+6}_{-16} (55^{+9}_{-22})$	$36^{+10}_{-15} (45^{+13}_{-19})$	$2.5^{+2.1}_{-0.7} (1.8^{+1.7}_{-0.5})$	$13^{+55}_{-8} (6^{+28}_{-4})$
GW190602_175927	$42^{+7}_{-17} (48^{+14}_{-19})$	$31^{+13}_{-16} (33^{+15}_{-16})$	$1.4^{+1.5}_{-0.5} (1.1^{+1.4}_{-0.4})$	$10^{+65}_{-7} (6^{+46}_{-4})$
GW190706_222641	$39^{+10}_{-15} (42^{+17}_{-17})$	$29^{+12}_{-13} (29^{+13}_{-13})$	$1.7^{+1.8}_{-0.5} (1.6^{+1.7}_{-0.6})$	$5^{+26}_{-3} (4^{+22}_{-2})$

Under the hypothesis that the listed events are lensed signals from intrinsically lower-mass binary populations with $\mu > 2$, this table lists the favored source masses, redshifts, and magnifications for the BNS and neutron star–black hole (NSBH) (top) and BBH (bottom) high-mass events. For the BBHs, two sets of numbers are given for different assumptions about the edge of the pair instability supernova (PISN) mass gap (a cut at $50 M_\odot$ and $65 M_\odot$). For the BNSs, we presume that they originate from the Galactic BNS population. To interpret the heavy BBHs as lensed signals originating from the assumed lower-mass population, they should be magnified at a moderate magnification $\mu \sim \mathcal{O}(10)$ at $z \sim 1$ to 2. The BNS and NSBH events would require extreme magnifications.

magnification. Therefore, we can compute the magnification posterior and other parameters by simply re-weighting existing posteriors.

Studies along these lines were already done for the GW190425 BNS event by Pang et al. [288] and for the GW190521 BBH event in Abbott et al. [18]. Here we extend the approach to cover additional interesting O3a events, focusing on two cases: (i) the (apparently) most massive observed BBHs, and (ii) sources with an (apparent) heavy neutron star component. In the BBH case, we take the prior over component masses, m_1 and m_2 , and redshift, z of the source $p(m_1, m_2, z)$ from the power-law BBH population model used in Abbott et al. [10] for O1 and O2 observations, with a mass power-law index $\alpha = 1$, mass ratio power-law index $\beta_q = 0$, and minimum component mass $m_{\min} = 5 M_\odot$, and assume an absence of BBHs above the pair instability supernova (PISN) mass gap. As in the previous GW190521 study [18], we consider two different values to account for uncertainties on the edge of the PISN gap, $m_{\max} = (50, 65) M_\odot$. Such a simple model is adequate for this analysis because our analysis results are most sensitive to the mass cut (highest masses allowed by the prior) and less sensitive to the specific shape of the mass distribution. For events

with an apparent heavy neutron star component, we assume a Galactic BNS prior following a total mass with a $2.69 M_\odot$ mean and $0.12 M_\odot$ standard deviation [160]. In both cases, the magnification could explain the apparent high mass of the events from the LIGO–Virgo observations.

We assume that the redshift prior $p(z) \propto \tau(z)dV_c/dz$, where the optical depth of lensing by galaxies or galaxy clusters $\tau(z) \propto D_c(z)^3$ [188]. The redshift dependence of the optical depth is approximately the same for both galaxies and galaxy clusters, while the overall scaling can change [168]. We use the lensing prior $p(\mu|\mathcal{H}_{\text{SL}}) \propto \mu^{-3}$ [64] with a lower limit $\mu > 2$ appropriate to strong lensing [274]. This prior is appropriate when we are in the high-magnification, strong lensing limit, i.e. assuming that the observed masses are highly biased. We do not consider weak lensing, which does not produce multiple images and would require expanded future GW data sets to study [264, 265].

We analyze all O3a BBH events with the primary mass above $50 M_\odot$ at 90% probability using the Bayesian inference posterior samples released with GWTC-2 [21, 181]. Moreover, we analyze GW190425, a high-mass BNS [12], and GW190426_152155, a low-significance potential neutron star–black hole (NSBH) event [21] which was investigated as a possible lensed BNS event [322]. We use the results for the IMRPHENOMPv2 waveform [68, 184] for most of the events. For GW190521, where higher-order multipole moments are important to include in the analysis [18], we adopt the NRSUR7DQ4 waveform [354] results as in Abbott et al. [16]. Furthermore, for GW190425 [12], we use the IMRPHENOMPv2_NRTIDAL [140] low-spin samples. Results are summarized in Table 8.4.1.

To interpret the heavy BBHs as lensed signals originating from the assumed lower-mass population, they should be magnified at a moderate magnification $\mu \sim 10$ at $z \sim 1 - 2$. Depending on the lens model, this magnification may imply a moderate chance of an observable multi-image counterpart as events closer to the caustic curves experience more substantial magnifications. Consequently, they often produce events with similar magnification ratios and shorter time delays (comparable magnifications and shorter time delays can be derived from the lens’s symmetry, although if lensing by substructures or microlenses is present, the magnifications between images can differ even in the high-magnification limit). However, we could not identify any multi-image counterparts for any of the high-mass events in our multiple image search (Sec. 8.5).

The BNS and NSBH events, on the other hand, would require extreme magnifi-

cations (68^{+163}_{-44} and 497^{+452}_{-272} , respectively) to be consistent with the Galactic BNS distribution. At these magnifications, we would expect the source to be close to a caustic, and therefore it may be possible that the presence of microlenses would produce observable effects [138, 139, 256, 286]. Moreover, the event would likely be multiply imaged [64, 281]. A more detailed follow-up study to quantify the likelihood of multiple images and microlensing could produce more stringent evidence for the lensing hypothesis for these events. We will briefly comment on these events in the context of multi-image and microlensing results in the sections that follow.

At this stage, we cannot set robust constraints on the lensing hypothesis based on the magnification alone. Moreover, as detailed in the following section, we have also not found any other clear evidence to indicate that these GW events are lensed. The prior lensing rate disfavors the lensing hypothesis for most standard binary population and lens models, as discussed in Sec. 8.3. However, if other BBH formation channels exist that produce an extensive number of mergers at high redshift, the lensing rates can change. In the future, more quantitative constraints could be set by connecting the inferred magnifications with lens modeling to make predictions for the appearance of multiple images or microlensing effects.

8.5 Search for multiple images

In addition to magnification, strong lensing can produce multiple images of a single astrophysical event. These multiple images appear at the GW detectors as repeated events. The images will differ in their arrival time and amplitude [186, 188, 223, 247, 361]. The sky location is the same within the localization accuracy of GW detectors, given that the typical angular separations are of the order of arcseconds. Additionally, lensing can invert or Hilbert transform the image [119, 154], introducing a frequency-independent phase shift. This transformation depends on the image type, set by the lensing time delay at the image position: Type-I, II, and III correspond to a time-delay minimum, saddle point, and maximum, respectively [154].

The multiply imaged waveforms $\{\tilde{h}_j^L\}$ of a single signal \tilde{h} then satisfy [119, 154]

$$\tilde{h}_j^L(f; \theta, \mu_j, \Delta t_j, \Delta \phi_j) = \sqrt{|\mu_j|} \tilde{h}(f; \theta, \Delta t_j) \exp(i \text{sign}(f) \Delta \phi_j), \quad (8.5)$$

where $\sqrt{|\mu_j|}$ is the lensing magnification experienced by the image j and $\Delta \phi_j = -\pi n_j/2$ is the Morse phase, with index $n_j = 0, 1, 2$ for Type-I, II, and III images. $\tilde{h}(f; \theta, \Delta t_j)$ is the original (unlensed) waveform before lensing, but evaluated as arriving with a time delay Δt_j . The multi-image hypothesis then states that most

parameters measured from the different lensed images of the same event are consistent.

The relative importance of different parameters for the overall consistency under the multi-image hypothesis will vary for different events. For example, the sky localization match will have greater relevance for well-localized, high-SNR events. Similarly, the overlap in measured chirp mass $(1+z)\mathcal{M} = (1+z)(m_1 m_2)^{3/5}/(m_1 + m_2)^{1/5}$ will be more significant when the uncertainty in that parameter is lower, although in this case the underlying astrophysical mass distribution will play a key role. The similarities in other parameters such as mass ratios or spins will be more important when they depart from the more common astrophysical expectations. Evidence of strong lensing could also be acquired with a single Type-II (saddle point) image if the induced waveform distortions in the presence of higher modes, precession, or eccentricity are observed [154]. Such evidence is unlikely to be observed without next-generation detectors [360].

In this section, we perform three distinct but related analyses. First, we test the lensed multi-image hypothesis by analyzing, for all pairs of O3a events from GWTC-2, the overlap of posterior distributions previously inferred for the individual events. This allows us to set ranking statistics to identify an initial set of candidates for lensed multiple images. We perform a more detailed joint-PE analysis for these most promising pairs, considering all potential correlations in the full parameter space and the image type. This joint analysis provides a more solid determination of the lensing probability for a given GW pair. Finally, we search for additional sub-threshold candidates that could be multiply imaged counterparts to the previously considered events: some counterpart images can have lower relative magnification compared with the primary image and/or fall in times of worse detector sensitivity or antenna patterns, and hence may not have passed the detection threshold of the original broad searches. According to the predictions of the expected lensing time delays and the rate of galaxy and galaxy cluster lensing [122, 281, 324], we expect it to be less likely for counterpart images to the O3a events to be detected in observing runs O1 or O2. Relative lensing rates for galaxies and clusters are given in Table 8.3.1. Thus, we only search for multiple images within O3a itself.

Previous studies have also searched for multiple images in the O1–O2 catalog GWTC-1 [76, 122, 186, 223, 233, 247]. The first search for GW lensing signatures in O1 and O2 focused on the posterior overlap of the masses, spins, binary orientation and sky positions [186] and the consistency of time delays with expectations for

galaxy lenses, but found no conclusive evidence of lensing. The search did uncover a candidate pair GW170104–GW170814 with a relatively high Bayes factor of $\gtrsim 200$. Still, this study disfavored the candidate due to its long time delay and the low prior probability of lensing. In parallel, Broadhurst et al. [76] suggested that the candidate pair GW170809–GW170814 could be lensed, but this claim is disfavored by more comprehensive analyses [186, 233]. Both Li et al. [223] and McIsaac et al. [247] performed searches for sub-threshold counterparts to the GWTC-1 events, identifying some marginal candidates but finding no conclusive evidence of lensing. More recently, Dai et al. [122] and Liu et al. [233] searched for lensed GW signals including the analysis of the lensing image type, which can be described through the Morse phases, $\Delta\phi_j$ in Eq. (8.5). These analyses have revisited the pair GW170104–GW170814 and demonstrated that the Morse phase is consistent with the lensed expectation but would require Type-III (time-delay maximum) images, which are rare from an observational standpoint. Dai et al. [122] also pointed out that a sub-threshold trigger, designated by them as GWC170620, is also consistent with coming from the same source. However, the required number and type of images for this lens system make the interpretation unlikely given current astrophysical expectations. Also, two same-day O3a event pairs (on 2019 May 21 and 2019 August 28) have already been considered elsewhere, but were both ruled out due to vanishing localization overlap [18, 320].

Posterior-overlap analysis

As a consequence of degeneracies in the measurements of parameters, the lensing magnification can be absorbed into the luminosity distance (Sec. 8.4), the time delay can be absorbed into the time of coalescence, and, when the radiation is dominated by $\ell = |m| = 2$ multipole moments, the phase shifts introduced by lensing (the Morse phases) can be absorbed into the phase of coalescence. The multi-image hypothesis then states that all other parameters except the arrival time, luminosity distance, and coalescence phase are the same between lensed events, and thus there should be extensive overlap in their posterior distributions, even if those have been inferred without taking lensing into account.

Therefore, we use the consistency of GW signals detected by LIGO and Virgo to identify potential lensed pairs. Following [188], we define a ranking statistic $\mathcal{B}^{\text{overlap}}$ to distinguish candidate lensed pairs from unrelated signals:

$$\mathcal{B}^{\text{overlap}} = \int d\Theta \frac{p(\Theta|d_1) p(\Theta|d_2)}{p(\Theta)}, \quad (8.6)$$

where the parameters Θ include the redshifted masses $(1+z)m_{1,2}$, the dimensionless spin magnitudes $\chi_{1,2}$, the cosine of spin tilt angles $\theta_{1,2}$, the sky location $(\alpha, \sin \delta)$, and the cosine of orbital inclination θ_{JN} , but they do not include the full 15-dimensional set of parameters Θ to ensure the accuracy of the kernel density estimators (KDEs) that we use to approximate the posterior distributions $p(\Theta|d_{1,2})$ for each event when evaluating Eq. (8.6). Here, $p(\Theta)$ denotes the prior on Θ .

The accuracy of the KDE approximation was demonstrated in Haris et al. [188] through receiver operating characteristic curves with simulated lensed and unlensed BBH events. To improve the accuracy further, we compute the sky localization (α, δ) overlap separately from other parameters and combine it with the overlap from the remaining parameters. Splitting the two overlap computations is justified because the posterior correlations of (α, δ) with other parameters are minimal.

We use posterior samples [181] obtained using the LALINFERENCE software package [356] with the IMRPHENOMPv2 waveform model [104, 184] for most of the events. However, for GW190521, we use NRSUR7DQ4 [354] posteriors, and for GW190412 and GW190814 we use IMRPHENOMPv3HM [210] posteriors. The prior $p(\Theta)$ is chosen to be uniform in all parameters. The component mass priors have the bound $(2\text{--}200 M_\odot)$. Equation (8.6) then quantifies how consistent a given event pair is with being lensed. In our analysis, we omit the BNS event GW190425 [12] because it was detected at relatively low redshift, and hence we expect the probability of it being lensed to be very small.

In addition to the consistency of the frequency profile of the signals (as measured by the posterior overlap), the expected time delays Δt between lensed images follow a different distribution than for pairs of unrelated events. Following Haris et al. [188], we define

$$\mathcal{R}^{\text{gal}} = \frac{p(\Delta t|\mathcal{H}_{\text{SL}})}{p(\Delta t|\mathcal{H}_{\text{U}})}, \quad (8.7)$$

where $p(\Delta t|\mathcal{H}_{\text{SL}})$ and $p(\Delta t|\mathcal{H}_{\text{U}})$ are the prior probabilities of the time delay Δt under the strongly lensed and unlensed hypotheses, respectively. Here $p(\Delta t|\mathcal{H}_{\text{U}})$ is obtained by assuming that the GW events follow a Poisson process. We use a numerical fit to the time-delay distribution $p(\Delta t|\mathcal{H}_{\text{SL}})$ obtained in Sec. 8.3 for the SIS galaxy lens model, with a merger rate density given by R_{min} in Eq. (8.10). Equation (8.7) provides another ranking statistic to test the lensing hypothesis, based on the time delay, though subject to some astrophysical uncertainties (see discussion in Sec. 8.3). The time-delay distribution does not include galaxy cluster lenses, which may be responsible for long time delays of several months or more. We also do

not model detector downtime, but we expect the different contributions to the time delay to average out across a longer time period.




Figure 8.5.1: Scatter plot of the ranking statistics $\log_{10} \mathcal{B}^{\text{overlap}}$ and $\log_{10} \mathcal{R}^{\text{gal}}$ for a subset of event pairs that have both $\mathcal{B}^{\text{overlap}} > 50$ and $\mathcal{R}^{\text{gal}} > 0.01$. The dashed lines denote the significance levels of the combined ranking statistics (in terms of Gaussian standard deviations), obtained by simulating unlensed event pairs in Gaussian noise matching the O3a sensitivity of the LIGO–Virgo network. We identify several high $\mathcal{B}^{\text{overlap}} > 50$ candidates, which we follow up on with a detailed joint-PE analysis. We have used abbreviated event names, quoting the last 4 digits of the date identifier (see Table 8.5.1 for full names).

To estimate the significance of the combined ranking statistic, $\log_{10}(\mathcal{B}^{\text{overlap}} \times \mathcal{R}^{\text{gal}})$ computed for O3a event pairs, we perform an injection campaign. For the injection campaign, we sample component masses $m_{1,2}$ from a power-law distribution [4] in the range $(10\text{--}50 M_{\odot})$. We assume that the redshift distribution follows population synthesis simulations of isolated binary evolution [56, 57, 65, 143, 148, 244, 273, 311]; in particular, for illustration purposes, we show results using the redshift

evolution from Belczynski et al. [55, 58], but for the local universe that we look at ($z < 2$), other models produce qualitatively similar results. All other parameters are sampled from uninformative prior distributions [188]. We inject the simulated signals into Gaussian noise with O3a representative spectra for a LIGO–Virgo detector network. We compute $\mathcal{B}^{\text{overlap}}$ and \mathcal{R}^{gal} for all possible pairs in the injection set to obtain the false-alarm probability for one pair $\text{FAP}^{\text{pair}}(x)$ at different levels x of combined statistics by counting the number of simulated pairs with $\log_{10}(\mathcal{B}^{\text{overlap}} \times \mathcal{R}^{\text{gal}}) > x$. Then the probability of at least one of the N event pairs in GWTC-2 to cross the threshold can be estimated as $\text{FAP}^{\text{cat}}(x) = 1 - [\text{FAP}^{\text{pair}}(x)]^N$. We then obtain the σ levels of significance shown in Fig. 8.5.1 by assuming $\text{FAP}^{\text{cat}}(x)$ follows the complementary error function.

In Fig. 8.5.1 we show the scatter plot of $\log_{10} \mathcal{B}^{\text{overlap}}$ and $\log_{10} \mathcal{R}^{\text{gal}}$ for the O3a event pairs that have high combined ranking statistic. The dashed lines represent different significance levels as obtained from the simulations. The event pair GW190728_064510–GW190930_133541 gives the highest combined ranking statistic, $\log_{10}(\mathcal{B}^{\text{overlap}} \times \mathcal{R}^{\text{gal}}) = 3.6$; however, as can be seen from Fig. 8.5.1, its significance is above 1σ (68%) but much below the 2σ (95%) significance level.

To follow up on the most promising event pairs with the more detailed joint-PE analysis in the next section, we make a selection based on just the posterior overlap ranking statistic, $\mathcal{B}^{\text{overlap}}$, rather than the combined ranking statistic, $\mathcal{B}^{\text{overlap}} \times \mathcal{R}^{\text{gal}}$, because \mathcal{R}^{gal} depends strongly on the lens model. That is, we do not rule out any candidates based on \mathcal{R}^{gal} . Our aim in the next section is to understand the high $\mathcal{B}^{\text{overlap}}$ event pairs in greater detail without resorting to any specific lens model. We thus select the most promising event pairs from Fig. 8.5.1, i.e., those with $\mathcal{B}^{\text{overlap}} > 50$, and carry out the joint-PE analysis in the next section. The 19 selected pairs are listed in Table 8.5.1.

Joint parameter estimation analysis

Here we follow up on the most significant pairs of events from the posterior-overlap analysis with a more detailed but more computationally demanding joint-PE analysis. The benefit of this analysis is that it allows for more stringent constraints on the lensing hypothesis by investigating potential correlations in the full parameter space of BBH signals, instead of marginalizing over some parameters. Moreover, it also includes a test for the lensing image type by incorporating lensing phase information.

We perform our analysis using two independent pipelines, a LALInference-based

pipeline [233] and a BILBY-based pipeline [HANABI; 235], giving us additional confidence in our results. Unlike the posterior-overlap analysis, the joint-PE analysis does not start from existing posterior samples. Instead, we start the inference directly using the detector strain data. In both pipelines, we follow the same data selection choices (calibration version, available detectors for each event, and noise subtraction procedures) as in the original GWTC-2 analysis [21], with special noise mitigation steps (glitch subtraction and frequency range limitations) taken for some events as listed in Table V of that paper. However, the two pipelines use different waveform models. In this section, we first describe how we quantify the evidence for the strong lensing hypothesis, then detail the two pipelines and finally present the results.

The coherence ratio and the Bayes factor

There will be three types of outputs for the joint-PE analysis. First, we compute a coherence ratio C_U^L , which is the ratio of the lensed and unlensed evidences, neglecting selection effects and using default priors in the joint-PE inference. We treat this as a ranking statistic, which quantifies how consistent two signals are with the lensed hypothesis. Large coherence ratios indicate that the parameters of the GWs agree with the expectations of multiple lensed events. This occurs, for example, when the masses and sky localization coincide. However, the coherence ratio does not properly account for the possibility that the parameters overlap by chance.

The likelihood that GW parameters overlap by chance sensitively depends on the underlying population of sources and lenses. For example, if there existed formation channels that produced GWs with similar frequency evolutions (as expected of lensing), the likelihood of an unlensed event mimicking lensing would increase substantially. Thus, we introduce a second output, the population-weighted coherence ratio $C_U^L|_{\text{pop}}$, which incorporates prior information about the populations of BBHs and lenses. The value of $C_U^L|_{\text{pop}}$ is subject to the choice of both the BBH and lens models.

Similarly, the probability that two signals agree with the multiple-image hypothesis is altered through selection effects, as some masses and sky orientations are preferentially detected. Thus, we also include the selection effects, which gives us our final output, the Bayes factor \mathcal{B}_U^L . The \mathcal{B}_U^L quantifies the evidence of the strong lensing hypothesis for a given detector network and population model. For the full derivations and detailed discussion on the difference between the coherence ratio

and the Bayes factor, see [235].

LALInference-based pipeline

For the LALInference-based pipeline, we adopt the method presented by Liu et al. [233], which was first used for analyzing pairs of events from GWTC-1 [11]. LALInferenceNest [356] implements nested sampling [321], which can compute evidences without explicitly carrying out the high-dimensional integral while sampling the posteriors. The LALInference-based pipeline uses the IMRPhenomD waveform [193, 209], which is a phenomenological model that includes the inspiral, merger, and ringdown phases but assumes non-precessing binaries and only $\ell = |m| = 2$ multipole radiation. This is motivated by the fact that most events detected so far are well described by the dominant multipole moment [11, 21]. Higher-order multipole moments, precession, or eccentricity could lead to non-trivial changes to the waveform for Type-II images, but such waveforms cannot currently be used with this pipeline. For a discussion of the events within GWTC-2 displaying measurable higher-order multipole moments or precession, see Appendix A of Abbott et al. [21].

As in the posterior-overlap analysis, we expect observed, lensed GWs to share the same parameters for the redshifted masses, spins, sky position, polarization angle and inclination, $\{(1+z)m_1, (1+z)m_2, \chi_1, \chi_2, \alpha, \delta, \psi, \theta_{JN}\}$. Hence, we force these parameters to be identical under the lensing hypothesis. For the unlensed hypothesis, we sample independent sets of parameters for each event. This is equivalent to performing two separate nested sampling runs and then combining their evidence. In total, LALInference samples in an 11-dimensional parameter space and provides C_U^L as the output.

We sample the apparent luminosity distance of the first event D_L^1 and the relative magnification μ_r [361] instead of the luminosity distance of the second event D_L^2 , using the relation $\sqrt{\mu_r} = D_L^1/D_L^2$. Since our waveform only includes the dominant $\ell = |m| = 2$ multipole moments, the lensing Morse phase is modeled by discrete shifts in the coalescence phase ϕ_c by an integer multiple of $\pi/4$ [with relation to the lensing phase shift $\Delta\phi = 2\Delta\phi_c$, 119, 154]. Thus, we consider all possible relative shifts $\Delta\phi_c \in \{0, \pi/4, \pi/2, 3\pi/4\}$ between two GW signals.

We set a uniform prior in $\log[(1+z)m_1]$ and $\log[(1+z)m_2]$ for both the lensed and unlensed hypothesis. The minimum and maximum component masses are respectively $3 M_\odot$ and $330 M_\odot$, with a minimum mass ratio of $q = m_2/m_1 = 0.05$.

This choice reduces the prior volume by $10^2 - 10^3$ compared to the uniform prior used in GWTC-2 [see 233, for discussion]. For the other parameters, the prior for the luminosity distance is $p(D_L) \propto D_L^2$ up to 20 Gpc, while the spins are taken to be parallel to the dimensionless orbital angular momentum with a uniform prior on the z components between -0.99 (anti-aligned) and $+0.99$ (aligned).

The HANABI pipeline

The HANABI pipeline, on the other hand, adopts a hierarchical Bayesian framework that models the data generation process under the lensed and the unlensed hypothesis. This pipeline uses the IMRP_{HENOMXPHM} waveform [295], which models the full inspiral–merger–ringdown for generic precessing binaries including both the dominant and some sub-dominant multipole moments. Therefore, the parameter space of HANABI enlarges to 15 dimensions.

HANABI differs from the LALINFERENCE-based pipeline in the treatment of the Morse phase. Here the lensing phase is directly incorporated in the frequency-domain waveform, accounting for any possible distortion of Type-II images [120, 154, 235]. Moreover, the lensed probability is computed by considering all possible combinations of image types with a discrete uniform prior [235]. For this reason, HANABI only produces one evidence per pair, and not one for each discrete phase difference as the LALINFERENCE-based pipeline. Unlike the LALINFERENCE-based pipeline, HANABI samples the observed masses in a uniform distribution. The mass ranges are different for each event pair, but an overall reweighting is applied later (see below). The rest of the prior choices for the intrinsic parameters are the same as for the LALINFERENCE-based pipeline with the addition of a discrete uniform prior on the Morse phase and isotropic spin priors.

In addition to computing the joint-PE coherence ratio, HANABI also incorporates prior information about the lens and BBH populations, as well as selection effects. In particular, the BBH population is chosen to follow a POWER LAW + PEAK MODEL in the primary mass following the best-fit parameters in Abbott et al. [22]. Similarly, the secondary mass is fixed to a uniform distribution between the minimum and the primary mass. HANABI also uses an isotropic spin distribution and merger rate history following Model A in Sec. 8.3. The lens population is modeled by the optical depth described in Hannuksela et al. [186] and a magnification distribution $p(\mu) \propto \mu^{-3}$ for $\mu \geq 2$. HANABI is thus able to output C_U^L , $C_U^L|_{\text{pop}}$ and \mathcal{B}_U^L . However, HANABI does not include any preference for a particular type of image, i.e., HANABI

uses a discrete, uniform prior for the Morse phase shift $\Delta\phi_j$.

Table 8.5.1: Summary of joint-PE results for event pairs in O3a.

Event 1	Event 2	$\log \mathcal{R}^{\text{gal}}$	$\log_{10}(C_{\text{U}}^{\text{L}})$ LALInference ($\Delta\phi$: 0, $\pi/2$, π , $3\pi/2$)	$\log_{10}(C_{\text{U}}^{\text{L}} _{\text{pop}})$ HANABI	$\log_{10}(\mathcal{B}_{\text{U}}^{\text{L}})$ HANABI
GW190412	GW190708_232457	−1.6	(+1.0, −9.7, −22.8, −4.4)	−6.6	−9.7
GW190421_213856	GW190910_112807	−	(+ 4.5 , +2.5, −1.5, −0.0)	−0.7	− 3.8
GW190424_180648	GW190727_060333	−1.8	(+ 4.9 , +0.0, +1.1, + 4.0)	−0.8	− 3.9
GW190424_180648	GW190910_112807	−	(+2.5, + 4.7 , + 4.3 , +1.6)	−0.8	− 3.9
GW190513_205428	GW190630_185205	−0.6	(+0.8, + 4.3 , −1.9, −6.5)	−2.4	− 5.5
GW190706_222641	GW190719_215514	+0.4	(+2.4, +2.4, −0.0, −0.5)	−0.3	−3.4
GW190707_093326	GW190930_133541	−1.5	(−4.6, −4.3, −3.5, −4.1)	−9.4	−12.5
GW190719_215514	GW190915_235702	−0.9	(+3.5, −2.1, −0.1, + 4.1)	−0.7	− 3.8
GW190720_000836	GW190728_064510	+0.5	(−1.4, −0.9, −4.5, −5.4)	−6.7	−9.8
GW190720_000836	GW190930_133541	−1.2	(−3.5, −2.8, −3.9, −3.9)	−9.2	−12.3
GW190728_064510	GW190930_133541	−1.1	(−3.6, −2.5, −3.1, −2.9)	−8.5	−11.6
GW190413_052954	GW190424_180648	+0.4	(+0.6, −0.9, +0.4, −0.0)	−1.6	−4.7
GW190421_213856	GW190731_140936	−2.1	(+3.1, −1.9, +2.5, + 5.2)	−0.2	− 3.3
GW190424_180648	GW190521_074359	−0.1	(+1.3, +3.8, +3.7, + 4.4)	−2.0	− 5.1
GW190424_180648	GW190803_022701	−2.1	(+ 4.2 , +1.9, +2.6, +3.1)	−1.0	− 4.1
GW190727_060333	GW190910_112807	−0.6	(+1.8, +3.3, +3.7, +3.4)	−1.4	−4.5
GW190731_140936	GW190803_022701	+0.9	(+ 4.1 , +3.2, +2.2, +3.4)	−0.9	− 4.0
GW190731_140936	GW190910_112807	−0.5	(+0.1, + 4.5 , +0.8, −7.2)	−1.2	− 4.3
GW190803_022701	GW190910_112807	−0.4	(+ 4.0 , + 5.5 , + 4.7 , +2.6)	−0.1	− 3.2

We select those events with posterior overlap ranking statistic larger than 50. For each pair of events presented in the first two columns, the third column lists the time-delay ranking statistic \mathcal{R}^{gal} as described in Section 8.5. The next column gives the coherence ratio of the lensed/unlensed hypothesis C_{U}^{L} obtained with the LALInference-based pipeline, including the results for the four possible lensing phase difference $\Delta\phi = 2\Delta\phi_{\text{c}}$. We highlight in boldface those pairs with $\log_{10}(C_{\text{U}}^{\text{L}}) > 4$ for at least one Morse phase shift. The fifth and sixth columns correspond to the HANABI results for the population-weighted coherence ratio $C_{\text{U}}^{\text{L}}|_{\text{pop}}$ and the Bayes factor $\mathcal{B}_{\text{U}}^{\text{L}}$. All quantities are given in \log_{10} . All high coherence ratio events display a small Bayes factor when including the population priors and selection effects. For the pairs GW190421_213856–GW190910_112807 and GW190424_180648–GW190910_112807, the time delays between events are larger than what we expect for galaxy lenses in our simulation, and thus $\mathcal{R}^{\text{gal}} = 0$.

Results

Within the O3a events, the LALInference-based pipeline finds 11 pairs with $\log_{10}(C_U^L) > 4$, indicating high parameter consistency. We have checked that the results of the LALInference-based pipeline are qualitatively consistent with those from HANABI. This reinforces our previous argument that the shift in the coalescence phase is a good approximate description of the lensing Morse phase given that in the present catalog most events are dominated by the $\ell = |m| = 2$ multipole moments. However, because of the pair-dependent prior choices of HANABI, we do not present its raw C_U^L results in Table 8.5.1.

We then include our prior expectation on the properties of the lensed images (derived from our BBH and lens population priors) and selection effects when computing the population-weighted HANABI coherence ratio and the Bayes factors \mathcal{B}_U^L . The results are summarized in Table 8.5.1. The event pair GW190728_064510–GW190930_133541, which seemed the most promising from the overlap analysis in Sec. 8.5, is disfavored by both joint-PE pipelines. After the inclusion of the population prior and selection effects, none of the event pairs display a preference for the lens hypothesis ($\log_{10} \mathcal{B}_U^L < 0$).

The population-weighted coherence ratio and the Bayes factor are subject to the BBH and lens model specifications. The population properties are not inferred taking into account the possibility of lensing. This introduces an inevitable bias, but it can be justified *a posteriori* to be a good approximation given the expected low rate of strong lensing. Additionally, the population properties include significant uncertainties in the hyper-parameter estimates and presume a population model. In any case, to quantify this intrinsic uncertainty in the modeling, we consider different choices for the mass distribution and merger rate history. Varying the maximum BBH mass and the redshift evolution of the merger rate using the $R_{min}(z)$ and $R_{max}(z)$ of Model A in Sec. 8.3, we find that the strong lensing hypothesis is always disfavored. While these results are subject to assumptions on prior choices, our results are sufficient to reject the strong lensing hypothesis: Even if other prior choices favored the lensing hypothesis, the evidence would at best be inconclusive.

The impact of selection effects is considerable. Among other reasons, this is because present GW detectors preferentially observe higher mass events [164], making coincidences in observed masses more probable. Along the same lines, given the specific antenna patterns of the current network of detectors, GW events are preferentially seen in specific sky regions with characteristic elongated localization

areas [99], which favors the overlap between different events.

We also reanalyze the GW170104–GW170814 event pair in the O2 data previously studied by Dai et al. [122], Liu et al. [233]. Using the LALINFERENCE-based pipeline, Liu et al. [233] found that the coherence ratio, including selection effects associated with the Malmquist bias [243], is $\log_{10}(C_U^L) \approx 4.3$ for a $\pi/2$ coalescence phase shift. However, when including together population and selection effects with HANABI, we find that the evidence drastically reduces to a Bayes factor of $\log_{10}(\mathcal{B}_U^L) \approx -2.0$.

In addition to the Bayes factor, it is important to contrast the recovered number of candidate lensed pairs and their properties with astrophysical expectations. In Sec. 8.3 we found that the relative rate of GW events with at least two strongly lensed images above the detection threshold is below $\sim 1.3 \times 10^{-3}$ for all considered BBH population models. Thus, the lensing rate estimates significantly disfavor the lensing hypothesis *a priori*; even a moderate Bayes factor would not by itself yet make a compelling case for strong lensing. Additionally, the type of images, arrival times, and magnifications provide additional information on the lensing interpretation’s plausibility. For example, a quantification of the time-delay prior can be computed by multiplying the coherence ratio by \mathcal{R}^{gal} . However, our final conclusions do not depend on the prior information about the lensing time delays or the prior odds against lensing: the prior lensing knowledge further disfavors the strong lensing hypothesis, but we did not use it to rule out any candidates.

Although we do not find evidence of strong lensing, future electromagnetic follow-up of the candidates could allow for independent support for the hypothesis if we identified a lensed counterpart galaxy to these events [187, 302, 308, 317, 323–325, 374]. This identification could take place by matching GW and electromagnetic image properties when four GW images are available [187]. With two images, the number of hosts could also be constrained [317, 374], but to a lesser degree due to degeneracies with the lens and source alignment and uncertainties introduced by micro/millilensing – although strong lensing by galaxy clusters might allow us to identify a single cluster candidate [302, 308, 323–325]. Moreover, strong lensing could have produced additional images below the noise threshold. We perform a further investigation of such sub-threshold counterparts in the next section.

Search for sub-threshold lensed images

Here we search for sub-threshold counterpart images of the O3a events from GWTC-2 that would not have been identified as confident detections by the search pipelines

used in Abbott et al. [21]. As lensed images could in principle appear anywhere in the entire O3a data, we perform targeted template bank searches for these sub-threshold lensed counterparts over the whole O3a strain data set, following the data selection criteria described in Abbott et al. [21]. We employ two matched-filter searches based on the GstLAL [87, 183, 250, 309] and PyCBC [127, 277, 279, 351] pipelines, adapted to the lensing case in similar ways as in Li et al. [223] and McIsaac et al. [247].

Search methods and setups

The lensed hypothesis states that the intrinsic masses and spins will remain consistent between multiple lensed images of the same event. Hence, we can perform searches that specifically target sub-threshold lensed counterparts of known events by creating reduced banks of template waveforms with masses and spins close to those inferred for the primary event. We use the public posterior mass and spin samples released with GWTC-2 [181] to create these targeted template banks. This ensures that the known events will match well with the templates while simultaneously decreasing the FAR of the search for similar events, potentially returning new candidates that did not reach the search threshold in Abbott et al. [21]. GstLAL's reduced banks contain between 173 and 2698 templates per search, while for each PyCBC search we select a single aligned-spin template. The construction of these template banks closely follows Li et al. [223], McIsaac et al. [247] and is further detailed in Appendix 8.B. Template waveforms are generated using the aligned-spin SEOBNRv4_ROM waveform [69, 297, 298] for both pipelines and all events, with the exception of GW190425 in the PyCBC search, where we use the TaylorF2 model [47, 61–63, 66, 67, 80, 126, 161, 253, 258, 293, 312].

Given these template banks, each search pipeline proceeds with configurations and procedures as outlined in Abbott et al. [21] to produce a priority list of potential lensed candidates matching each target event. To rank these, each pipeline uses a different method to estimate FARs.

GstLAL first identifies matched-filter triggers from one or more of the Hanford, Livingston, and Virgo data streams. Coincidences are identified with the same settings as in Abbott et al. [21]. From each candidate's recovered parameters, a likelihood-ratio ranking statistic is computed [309]. Single-detector triggers are penalized using machine-learning based predictions [iDQ; 149, 176] whereas for coincident triggers, no data quality products are used. We estimate the FAR of

a trigger by comparing with the distribution of the ranking statistic from all non-coincident noise triggers, used to characterize the noise distribution, over the O3a data set.

PyCBC also first identifies single-detector matched-filter triggers, with a reduced clustering window compared to the GWTC-2 configuration (from 1 s to 0.01 s). These are tested for time coincidence between detectors and are required to have an $\text{SNR} \geq 4$ in at least two detectors. While in the GWTC-2 analysis the PyCBC search was limited to the Hanford and Livingston detectors, here we also include Virgo data, using the methods described in Davies et al. [127] to analyse the three detector network. FARs are estimated from a noise background measured using time-shifted data. All triggers within 0.1 s of the times of the events in GWTC-2 are removed from both the foreground (observed coincident events) and the background.

Candidates from both pipelines are further vetted by a sky localization consistency test against the targeted GWTC-2 event, as lensed images of the same event should come from consistent sky locations but the matched-filter searches do not check for this. For each new candidate, we generate a sky localization map $p(\Omega)$ using BAYESTAR [319], with Ω denoting parameters that specify the sky location. We compute the percentage overlap $O_{90\%CR}$ of the 90% credible regions between the sky localization $q(\Omega)$ of a GWTC-2 event and the sky localization $p(\Omega)$ of a sub-threshold event candidate as

$$O_{90\%CR} = \frac{\mathbf{1}_{90\%CR} [p(\Omega)q(\Omega)] d\Omega}{\min(\mathbf{1}_{90\%CR} [p(\Omega)] d\Omega, \mathbf{1}_{90\%CR} [q(\Omega)] d\Omega)}, \quad (8.8)$$

where $\mathbf{1}$ is the indicator function. To avoid false dismissal at this step, we only veto candidates with $O_{90\%CR} = 0$. All candidates with non-vanishing localization overlap are kept for further follow-up with data quality checks as discussed in Sec. 8.2 and with the joint-PE methods described in Sec. 8.5.

Results

In Table 8.5.2, we list the eight candidates with $\text{FAR} < 1$ in 16 years from the *individual* targeted searches for counterparts of the 39 detections reported in GWTC-2 found by at least one pipeline. Six of these are unique candidates. This number, compared with ~ 2 expected noise events above this FAR from the number of searches performed, is consistent with additional astrophysical signals being present in the data set. However, in this work, we do not assess in detail the probability of astrophysical origin for each of these. The reported FARs also do not indicate

how likely each trigger is to be a lensed counterpart of the targeted event, but only how likely it is to obtain a trigger with a similar ranking statistic from a pure noise background using these reduced template banks. Three of these candidates were also recovered with high probability of an astrophysical origin in the 3-OGC open-data search [278], which used a broad template bank. Five of them are also included with $p_{\text{astro}} > 0.5$ in the extended catalog GWTC-2.1 [33]. Candidates matching one or both of these catalogs are marked with footnotes in Table 8.5.2.

In contrast, Fig. 8.5.2 shows the *combined* search results from all 39 targets for each pipeline: GstLAL (top panel) and PyCBC (bottom panel), excluding triggers that correspond to other detections already reported in GWTC-2. Each panel shows the cumulative number of coincident triggers (observed) with inverse FARs greater than or equal to a given threshold value. For GstLAL, the combined results are obtained by a search over all O3a data using a combined template bank from the 39 targeted banks. For PyCBC, the FARs are obtained from the individual searches, but for triggers being found in several single-template searches, their inverse FARs are summed. In the same figure, we compare these results with estimated background distributions, accounting for the fact that we have re-analyzed the same data set of ~ 150 days multiple times, and find a slight excess in the rate of foreground triggers at high inverse FARs.

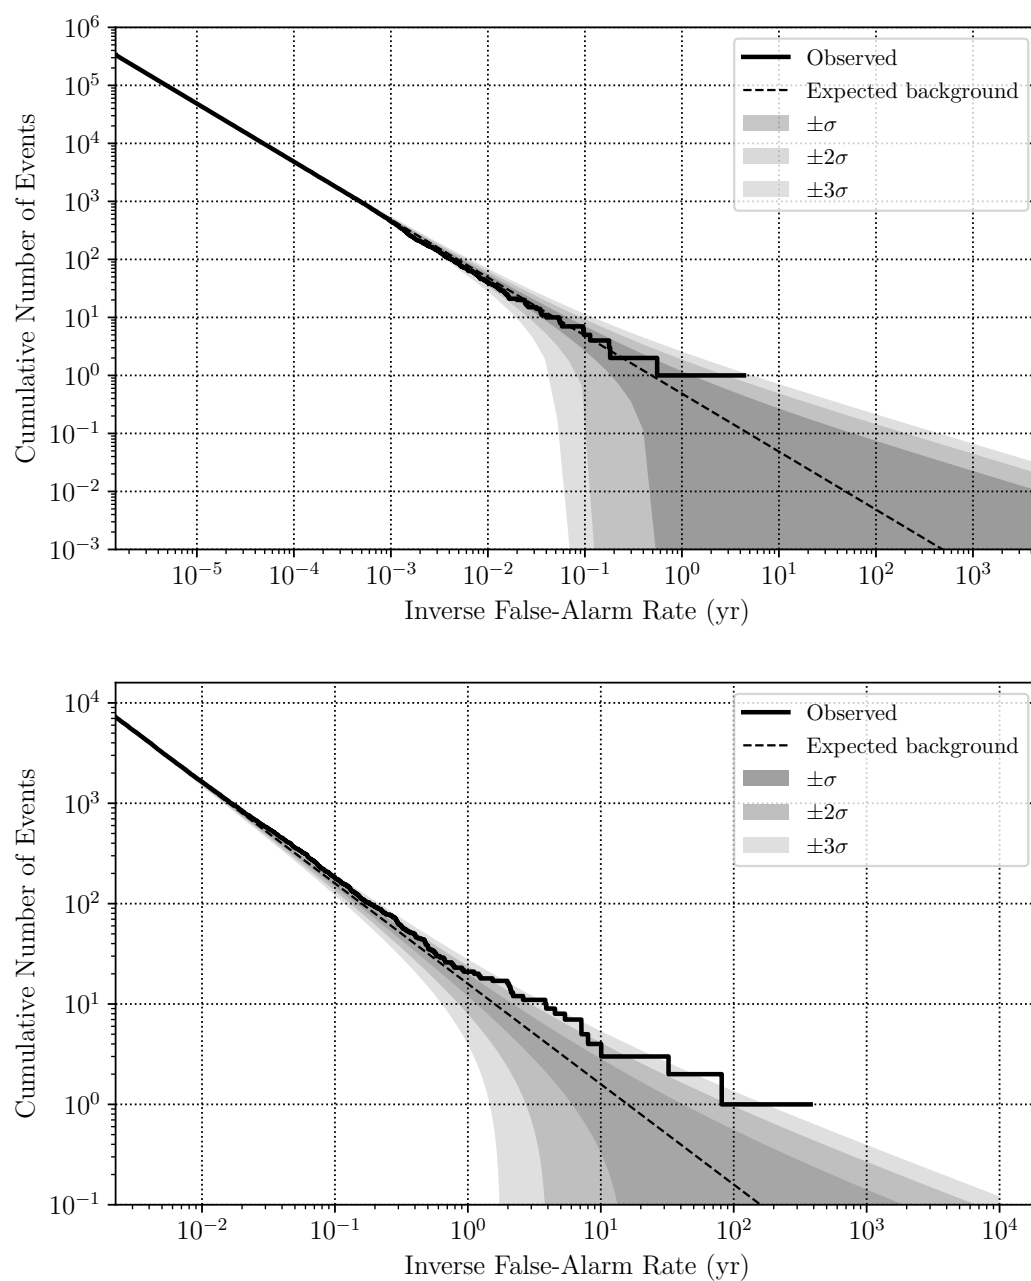


Figure 8.5.2: Combined results from the 39 sub-threshold searches with the GstLAL pipeline (top panel) and PyCBC pipeline (bottom panel). Each panel shows, as a solid line, the cumulative number of coincident triggers (observed) with inverse FARs greater than or equal to a given value. The dashed line is the expected distribution of background triggers, with the gray bands indicating uncertainties in multiples of the standard deviation σ of a Poisson distribution. For GstLAL, the results for this plot are obtained by a search over all O3a data using a combined bank from the 39 targeted banks. For PyCBC, the FARs are from the individual searches, but for triggers found by several of the single-template searches, their inverse FARs have been summed.

Table 8.5.2: Candidates from individual sub-threshold searches for strongly-lensed counterpart images of the 39 O3a events from GWTC-2.

UTC time	GWTC-2 target event	$ \Delta t $ [d]	FAR [yr ⁻¹]		$O_{90\%CR}$ [%]	$\log_{10} C_U^L$ (LALInference) ($\Delta\phi$: 0, $\pi/2$, π , $3\pi/2$)
			PyCBC	GstLAL		
190925_23:28:45 ^{a,b}	GW190828_06	28.69	0.003	98.681	0.0%	—
190426_19:06:42 ^b	GW190424_18	2.04	—	0.017	63.8%	(−5.8, −5.8, −5.9, −5.6)
190711_03:07:56	GW190421_21	80.23	0.032	0.341	1.2%	(+2.3, +1.1, +1.1, +2.6)
190725_17:47:28 ^{a,b}	GW190728_06	2.54	—	0.038	0.0%	—
190711_03:07:56	GW190731_14	20.46	0.045	0.944	2.9%	(+2.6, −1.2, −1.6, +0.9)
190805_21:11:37 ^b	GW190424_18	103.13	—	0.051	26.9%	(−1.1, +0.6, −0.3, −0.7)
190711_03:07:56	GW190909_11	60.36	0.053	1.196	12.6%	(+3.5, +2.2, +3.4, +2.9)
190916_20:06:58 ^{a,b}	GW190620_03	88.71	0.055	1.389	49.5%	(+1.7, +3.6, +2.1, −3.2)

The first column shows the UTC time of the newly found sub-threshold candidate (format YYMMDD_HHMMSS). The second column lists the targeted O3a event from the catalog GWTC-2 (Note that the event IDs have been shortened to only include up to “hour” of the UTC time of the events); see Table IV and Table VI of Abbott et al. [21] for details of these. The third column shows the absolute time difference between the candidate and the targeted event. The fourth and fifth columns show the corresponding FARs from the individual search for the target from the second column, from each of the two search pipelines (GstLAL and PyCBC), if the candidate has been recovered by it. The sixth column shows the percentage overlap of the 90% sky localization regions between the candidate and the targeted event, from the pipeline with the lower FAR. The seventh column shows the coherence ratio C_U^L for the pair from the LALInference joint-PE follow-up (only for candidate pairs with a localization overlap $> 0\%$). Candidates are only reported here if they pass a FAR threshold of < 1 in 16 years in at least one pipeline, and are sorted in ascending order by the lowest FAR from either pipeline. If the same new trigger was found with sufficient FAR by more than one search for different targets, all occurrences are included, and the PE follow-up is conducted separately for each pair. Candidates that have since also been reported by other searches are marked with footnotes.

[a] also included in 3-OGC [278]

[b] also included in GWTC-2.1 [33]

We perform follow-up analyses of the lensing hypothesis under the assumption of astrophysical origin, aiming to determine for each candidate pair in Table 8.5.2 whether it is more consistent with a pair of images of a single lensed event or with two independent astrophysical events. After taking into account the initial FAR thresholds, sky localization overlap, and data-quality checks, we have followed up six candidate pairs through LALInference joint Bayesian PE as described in Sec. 8.5. No special mitigation steps were required for data-quality reasons on any of the new candidates. The results are included in Table 8.5.2.

Compared with the results for GWTC-2 pairs in Table 8.5.1, the LALInference coherence ratios alone are insufficient to provide evidence of lensing while keeping in mind selection effects and prior odds. As another cross-check, we have also analysed the pair with the highest LALInference coherence ratio C_U^L (the candidate on 2019 September 16 found by the GW190620_030421 PyCBC search) with the HANABI pipeline described in Sec. 8.5. As with all pairs previously tested (see Table 8.5.1), after the inclusion of population priors and selection effects, there is no evidence favoring the lensing hypothesis for this pair either, with population-weighted coherence ratio $\log_{10}(C_U^L|_{\text{pop}}) = -0.1$ and Bayes factor $\log_{10}(\mathcal{B}_U^L) = -3.2$.

As lensing can produce more than two images of the same source, cases where several searches find the same trigger are of particular interest. We find that the same candidate on 2019 July 11 has been found with low FARs by three searches (targeting the GWTC-2 events GW190421_213856, GW190731_140936, and GW190909_114149). In addition, the trigger on 2019 August 05 is only found with sufficient FAR for inclusion in Table 8.5.2 by a single GstLAL search (for GW190424_180648), but was also recovered by those for GW190413_052954 and GW190803_022701 with FARs just below the cut. However, the GWTC-2 pairs involved in these possible quadruple sets have already been significantly disfavored by the HANABI analysis including population priors and selection effects. We also expect such multiple matches from an unlensed BBH population due to the clustering of the GWTC-2 events in parameter space [21, 22].

Also, as discussed in detail in McIsaac et al. [247], if any high-mass GW detections are interpreted as highly magnified images of lower-mass sources, then counterpart images for these would be more likely. However, we did not find any promising sub-threshold candidates for the five events discussed under the lensing magnification hypothesis in Sec. 8.4.

In summary, the sub-threshold searches can recover additional promising candidates

that were not included in GWTC-2, which match other events closely and, in that sense, are consistent with the lensing hypothesis. However, we do not find sufficient evidence that they are indeed lensed images, as the set of results is also consistent with a population of physically independent and only coincidentally similar events.

8.6 Search for microlensing effects

Microlensing by smaller lenses produces image separations on the order of microarcseconds. For GWs, it can also induce frequency-dependent wave-optics effects similar to femtolensing of light [270, 337]. More specifically, when the characteristic wavelengths are comparable to the Schwarzschild radius of the lens, i.e., $\lambda_{\text{GW}} \sim R_{\text{Sch}}^{\text{lens}}$, it causes frequency-dependent magnification of the waveform. Moreover, the characteristic lensing time-delay due to microlensed images can be shorter than the GW signal duration, causing potentially observable beating patterns on the waveform [90, 101, 105, 121, 138, 139, 205, 220, 256, 286], due to waveform superposition. To observe GW microlensing, we search for these beating patterns instead of the time-dependent change in the flux traditionally observed for microlensing in electromagnetic signals.

Here we search for microlensing by isolated point masses. The microlensed waveform has the form

$$h^{\text{ML}}(f; \theta_{\text{ML}}) = h^{\text{U}}(f; \theta) F(f; M_{\text{L}}^z, y), \quad (8.9)$$

where h^{ML} and h^{U} are the microlensed and unlensed waveforms in the frequency domain, respectively. θ represents the set of parameters defining an unlensed GW signal, while $\theta_{\text{ML}} = \{\theta, M_{\text{L}}^z, y\}$. $F(f; M_{\text{L}}^z, y)$ is the frequency-dependent lensing magnification factor, which is a function of the redshifted lens mass $M_{\text{L}}^z = M_{\text{L}}(1 + z_{\text{L}})$ and dimensionless impact parameter y , given in Eq. 2 of [220]. The search involves re-estimating the parameters of previously identified events under the microlensed hypothesis as defined in Eq. (9.2), including those of the potential lens.

To measure the evidence of lensing signatures in a signal, we define a Bayes factor $\mathcal{B}_{\text{U}}^{\text{ML}}$, which is the evidence ratio between the microlensed and unlensed hypotheses. Higher positive values correspond to support for lensing. [186] searched for similar beating patterns due to point mass lenses in the O1 and O2 data, using an upper lens mass prior cutoff $M_{\text{L}}^z \lesssim 10^5 M_{\odot}$. They reported no evidence for such lensing patterns above $\log_{10} \mathcal{B}_{\text{U}}^{\text{ML}} > 0.2$.

For O3a, we analyze the 36 events from [21] that confidently have both component masses above $3M_{\odot}$ and search for microlensing signatures following the same

method as in [186]. We perform PE using BILBY [49, 305] and the nested sampling algorithm dynesty [333]. For each event, we perform two PE runs using both unlensed and microlensed templates. For the unlensed case, which is similar to the usual PE analysis, equivalent prior settings and data dictionaries such as strain data and power spectral densities (PSDs) are used as in Abbott et al. [21]. The analysis uses the IMRPHENOMXPHM [295] waveform for most events, except for GW190521, which is analyzed using the NRSUR7DQ4 waveform [354] and for the least massive event GW190924_021846 where the IMRPHENOMPv2 waveform is used. The prior on M_L^z is log uniform in the range $[1-10^5 M_\odot]$, above which the effect of microlensing is relatively small for the LIGO–Virgo sensitivity band. The impact parameter prior is $p(y) \propto y$ between $[0.1, 3]$, chosen due to geometry and isotropy [220].

In Fig. 8.6.1 we show violin plots of marginalized posterior distributions for the redshifted lens mass for each event, as well as the Bayes factors between the microlensed and unlensed hypotheses. The broad M_L^z posteriors correspond to broad posteriors on the impact parameter y , which is not well constrained for unlensed cases. In terms of Bayes factors, there is no substantial evidence of microlensing with a maximum $\log_{10} \mathcal{B}_U^{\text{ML}} = 0.5$ for the event GW190910_112807. Additionally, as can be seen in Appendix 8.C, statistical fluctuations of the \log_{10} Bayes factors for injections without microlensing can be as high as 0.75. Thus, the observed Bayes factors are already by themselves consistent with random noise fluctuations and do not significantly favor the microlensing hypothesis for any of the events. The resulting posterior odds $\mathcal{O}_U^{\text{ML}}$, which are the products of Bayes factors and the low prior odds of microlensing [220], would be even lower. Thus, we find no evidence of microlensing in this study.




Figure 8.6.1: The marginalized posterior distribution of redshifted lens mass M_L^z and $\log_{10} \mathcal{B}_U^{\text{ML}}$ between microlensed and unlensed hypotheses. The corresponding \log_{10} Bayes factors are noted to the right of the plot. We find no evidence of microlensing by point mass lenses.

We searched for microlensing due to isolated point masses. More complex models in which point mass lenses embedded in an external macromodel potential such as galaxies and galaxy clusters [101, 138, 256] can produce additional modulation on the magnified waveform, which could also prove important in the LIGO–Virgo frequency band. Future searches could be extended to cover a broader range of microlensing models.

8.7 Conclusions and outlook

We have searched for gravitational lensing effects on the GW observations from O3a, the first half of the third LIGO–Virgo observing run, finding no strong evidence of lensing. First, we outlined estimates for the rate of strongly lensed GWs. Second, presuming a non-observation of lensing, we constrained the BBH merger-rate density at high redshift. Third, we used merger-rate density models obtained through the non-observation of a SGWB to estimate the GW lensing rate.

Next, we performed an analysis of apparent high-mass events under the hypothesis that they are lensed signals from lower-mass sources, finding that the highest-mass BBHs from O3a could be consistent with component masses below the PISN mass gap, while GW190425 and GW190426_152155 would require extreme magnifications to be compatible with the Galactic BNS population. This hypothesis is at the moment mainly disfavored by the expected lensing rates, but in the future, more quantitative constraints could also be set by connecting these magnification results with lens modeling to make predictions for the appearance of multiple images or the possibility of microlensing.

We then searched for signatures of multiple lensed images from a single source through several methods. We first investigated the parameter consistency among all pairs of O3a events from GWTC-2 using a posterior-overlap method, finding no significant event pairs but identifying several interesting candidates with high overlap.

We followed up on these candidate pairs using two detailed joint-PE analyses, finding high parameter consistency for 11 pairs. However, after the inclusion of a more appropriate population prior, selection effects, and the prior odds against the lensing hypothesis, these candidates do not provide sufficient evidence for a strong lensing claim.

Moreover, we used two targeted matched-filter approaches to search for additional lensed images of the known events that could be hidden beneath the thresholds of

the corresponding broader analyses used to produce GWTC-2, identifying six new candidates. After follow-up by joint PE, we found no evidence to conclude that any of these sub-threshold triggers are lensed images.

Finally, we analyzed 36 events from GWTC-2 for microlensing effects by performing full PE with waveforms incorporating microlensing by point mass lenses. We found no evidence of microlensing.

In summary, our results on O3a data are consistent with the expected low rate of lensing at current detector sensitivities. However, improved analysis methods and lens modeling may allow digging deeper into potential lensing effects. Electromagnetic follow-up of lensing candidates, even if they are not significant enough based on the GW data alone, could also be promising [187, 317, 324, 374]. With the current generation of detectors further improving their sensitivity and the global network being extended [9], the chances of detecting clear lensing signatures will improve, and the field will offer many possibilities at the latest with third-generation [14, 242, 296, 300] and space-based detectors [44, 190] and their expected cosmological reach.

APPENDIX

8.A Lensing statistics supplementary

Assuming a specific BBH formation channel, we can estimate the lensing rate for merger signals from that population. For example, suppose BBHs form as a consequence of isolated binary evolution. In that case, one can theoretically model BBH formation assuming that it traces the star-formation rate, modulated by the delay time distribution and by the stellar metallicity evolution [56, 57, 65, 143, 148, 244, 273, 311]. However, note that if the BBHs form through other means or through multiple channels, the merger-rate density could be different [e.g., 45, 46, 72, 132, 167, 254, 304, 367, 375].

Here we assume two models for the merger-rate density. We base the first model on the assumption that the merger-rate density of the observed BBHs traces the star-formation rate density and the BBHs originate from Population I/II stars.

In this work, we did not consider the contribution of Population III stars. Population III stars have not been observed yet, and their physical properties, binary fraction, and initial mass function are still a matter of debate [78, 194, 238, 240, 268, 280, 313]. As such, the contribution of Population III BBHs to gravitational-wave sources is also uncertain [e.g., 54, 71, 219, 232]. Should Population III stars dominated the BBH formation at high redshift, our results would need to be re-interpreted.

The first model, which we label Model A, uses the following fits that bracket the available population synthesis results from the literature [e.g., 56, 57, 65, 143, 148, 244, 273, 311]:

$$\begin{aligned}\mathcal{R}_m^{\min}(z_m) &= \frac{a_1 e^{a_2 z_m}}{a_3 + e^{a_4 z_m}} \text{Gpc}^{-3} \text{yr}^{-1}, \\ \mathcal{R}_m^{\max}(z_m) &= \frac{b_1 e^{b_2 z_m}}{b_3 + e^{b_4 z_m}} \text{Gpc}^{-3} \text{yr}^{-1},\end{aligned}\tag{8.10}$$

where the fitting parameters $a_1 = 58.497$, $a_2 = 2.06424$, $a_3 = 2.82338$, $a_4 = 2.52898$, $b_1 = 105356$, $b_2 = 1.30278$, $b_3 = 2714.36$, and $b_4 = 2.22903$.

We base the second model, Model B, on the assumption that the merger-rate density follows the Madau and Dickinson [239] ansatz:

$$\mathcal{R}_m(z_m; R_0, \alpha) = R_0 \frac{(1 + z_m)^\kappa}{1 + [(1 + z_m)/(1 + z_p)]^{(\gamma + \kappa)}}.\tag{8.11}$$

To constrain the merger-rate density at high redshift, we assume that no strong lensing has occurred (Sec. 9.5). We further assume that events occur following a Poisson process.

Let us now assume Model B for the merger-rate density, Eq. 8.11. The distribution of merger-rate density parameters, given that no strong lensing has occurred,

$$p(R_0, \kappa, \gamma, z_p | N, \{d_i\}) \propto \mathcal{W} \times p(R_0, \kappa, \gamma, z_p | \{d_i\}), \quad (8.12)$$

where $p(R_0, \kappa, \gamma, z_p | \{d_i\})$ follows the posterior distribution of parameters inferred from LIGO–Virgo population studies [22], and

$$\mathcal{W} = \frac{N_{\text{avg}}(R_0, \kappa, \gamma, z_p)^N \exp[-N_{\text{avg}}(R_0, \kappa, \gamma, z_p)]}{N!}, \quad (8.13)$$

with N being the number of observed, strongly lensed GW signals, and $N_{\text{avg}}(R_0, \kappa, \gamma, z_p)$ the expected number of events within a time Δt . Here, like in Sec. 8.3, we do not account for detector downtime, and instead as a proxy presume that the detectors are always online. The R_0 and κ value is measured at a low redshift [22]. The γ and z_p values are unconstrained here and thus match an uninformative prior, with $p(\gamma) = \mathcal{SN}(5, 10, 3)$ being a split normal distribution and $p(z_p)$ being uniformly distributed between $[0, 4]$. The above equations give all the necessary ingredients to forecast the rate of strongly lensed events and place constraints on the merger-rate density based on the number of lensed signals observed by LIGO and Virgo.

8.B Construction of sub-threshold counterpart search template banks

For the GstLAL and PyCBC searches for sub-threshold lensed counterparts (Sec. 8.5) the targeted template banks for each event are constructed starting from a certain choice of posterior distributions released with GWTC-2 [21, 181], aiming for a reduced-size template bank that is effective at recovering signals similar to the primary observed event.

For the GstLAL pipeline, we start, for all but three of the O3a events from GWTC-2, from non-spinning posteriors obtained with the IMRPHENOMD waveform [193, 209]. In three cases, we instead start from posteriors obtained with the IMRPHENOMPv2 waveform [68, 184], which includes spin precession. These exceptions are GW190413_052954, GW190426_152155, and GW190909_114149. We then choose subsets of the original broad template bank from the GWTC-2 analysis by comparing against the posteriors of each event, using the following steps as introduced by Li et al. [223]: We first draw $O(1000)$ of each event’s posterior samples with the highest likelihoods to account for the uncertainty in the

event’s measured mass and spin parameters. For each sample we simulate, using the aligned-spin SEOBNRv4_ROM waveform model [69, 297, 298], one signal with the event’s original optimal signal-to-noise ratio ρ_{opt} as given by Eq. (2) in [223] and nine extra signals with smaller ρ_{opt} , scaled by changing their effective distances D_{eff} [43]. The reduced template bank for an event is then constructed by searching the simulated data with the original GWTC-2 template bank (which also consists of SEOBNRv4_ROM waveforms) and keeping those templates that recover any of the simulated signals with a FAR < 1 in 30 days.

For PyCBC we select a single template for each search, choosing the maximum-posterior redshifted masses and aligned-spin components $\{(1+z)m_1, (1+z)m_2, \chi_1, \chi_2\}$ as estimated from a four-dimensional Gaussian KDE fit to the posterior samples from GWOSC [181] for these parameters. Where available, we use aligned-spin posterior samples. In the case of GW190412 and GW190814, we use samples generated using the SEOBNRv4_ROM waveform; for GW190426_152155 we use a mixture of samples generated using the SEOBNRv4_ROM_NRTidalv2_NSBH and IMRPhenomNSBH waveforms; and for GW190425 we use samples generated using the IMRPhenomD_NRTidal, TEOBResumS, and SEOBNRv4T_surrogate waveforms. If aligned-spin posteriors are not available in the GWOSC [181] data release, we use precessing posterior samples and marginalise over the transverse-spin components before applying the KDE. This produces an aligned-spin template with high matches at the peak of the posterior. In the case of GW190521, we use samples generated using the IMRPhenomPv3HM [210], NRSUR7DQ4 [354] and SEOBNRv4PHM [285] waveforms. For all other events, we use samples generated using the SEOBNRv4P and IMRPhenomPv2 waveforms.

These choices of waveforms and posterior samples are not necessarily optimal, but they are valid for this analysis in the sense that the recovery of similar waveforms with parameters close to the best-fit ones for the targeted GWTC-2 events has been verified through injection studies. In addition, in the actual searches, the targeted banks constructed in this way successfully recovered the corresponding GWTC-2 events in all GstLAL searches, while for PyCBC triggers within 0.1 s of the target events were excluded from the final trigger list, but in all cases where the original events were observed with two or more detector, a coincident trigger was also recovered in the targeted search. In future work, revisiting the choice of posterior samples used to construct template banks may further improve sub-threshold searches’ effectiveness.

8.C Injection study for microlensing analysis




Figure 8.C.1: Distribution of microlensing Bayes factors $\log_{10} \mathcal{B}_U^{\text{ML}}$ for unlensed simulated signals, recovered using a lensed template.

A high Bayes factor $\mathcal{B}_U^{\text{ML}}$ itself is not conclusive evidence of microlensing in an observed event. We have performed an injection study to explore the impact of statistical fluctuations on the Bayes factor obtained from unlensed signals. We generate unlensed injections by randomly drawing from the parameter space of precessing BBH systems. Simulated Gaussian noise is used considering nominal O3 sensitivity [9], and we use the IMRPHENOMPv2 waveform model [68, 184] for all simulated injections. The statistical fluctuations of $\log_{10} \mathcal{B}_U^{\text{ML}}$ for 100 unlensed injections recovered using lensed templates can be seen in Fig. 8.C.1 which shows that the typical values found are $\log_{10} \mathcal{B}_U^{\text{ML}} < 0.75$.

Chapter 9

SEARCH FOR LENSING SIGNATURES IN LVK’S FULL THIRD OBSERVING RUN

Note: This chapter is an adaptation of the LIGO-Virgo-KAGRA collaboration-wide publication:

LIGO and Virgo Scientific Collaboration, R. Abbott, ...(850 authors)..., **Alvin K. Y. Li**, ...(842 authors)..., “Search for gravitational-lensing signatures in the full third observing run of the LIGO-Virgo network”, *arXiv e-prints*, 2023.
DOI:10.48550/arXiv.2304.08393.

Alvin K. Y. Li is one of the editorial leads for writing this LIGO-Virgo-KAGRA collaboration-wide paper and led the analysis for the search for sub-threshold lensed gravitational waves (Section 3.1). The other related sections in this chapter are contributions from co-authors.

9.1 Introduction

In this chapter, we search for a variety of lensing signatures in the third LIGO Scientific, Virgo, and KAGRA (LVK) Collaboration Gravitational-Wave Transient Catalog (GWTC-3) [30] and study its implications for gravitational wave lensing. In particular, we expand on the lensing results presented for the first half of the third observing run of the LIGO–Virgo network (O3a) [23] by including the signals found in the second half of the third observing run (O3b) and by including additional analyses to further test the lensing hypothesis and interpret their outcomes. First, we search for the effects of strong lensing by studying the similarity and lensing evidence for pairs of binary black holes mergers. We consider both pairs of detected mergers (super-threshold) and pairs formed by detected mergers and candidates that nominally fall below the detection threshold (sub-threshold) with consistent waveform morphologies. Second, we search for evidence of microlensing induced by point-mass lenses. Finally, we constrain the expected rate of lensed signals, black holes merger-rate density, and the fraction of dark matter composed of compact objects.

It is important to note that GWTC-3 is a cumulative catalog describing all the

gravitational wave transients found in all observing runs to date: O1, O2, O3a, and O3b. O1 made observations between 2015 September 12 00:00 UTC to 2016 January 19 16:00 UTC, O2 between 2016 November 30 16:00 UTC to 2017 August 25 22:00 UTC, O3a between 2019 April 1 15:00 UTC to 2019 October 1 15:00 UTC, and O3b between 2019 November 1 15:00 UTC to 2020 March 27 17:00 UTC.

Results of all analyses in this paper and associated data products can be found in LIGO Scientific Collaboration and Virgo Collaboration [228]. gravitational wave strain data [182] and posterior samples [27] for all events from GWTC-3 are available from the Zenodo platform or the Gravitational Wave Open Science Center [35].

9.2 Data and Events

The analyses presented here expand on the lensing results from the first half of O3 (also referred to as O3a) by documenting new results from the second half of O3 (also referred to as O3b) using GWTC-3 [30]. The O3a lensing results paper [23] used the GWTC-2 catalog [21]. Since then, GWTC-2.1 [33] has reclassified 2 of the candidates used in the O3a lensing paper as having a probability of astrophysical origin of less than 0.5 and are not included in the results described here, specifically GW190424_180648 and GW190909_114149. GWTC-3 also includes 5 events that were identified by the O3a lensing sub-threshold search using the TESLA pipeline introduced in Chapter 7, namely GW190925_233845, GW190426_190642, GW190725_184728, GW190805_211137, and GW190916_200658.

Various instrumental upgrades have led to more sensitive data in O3b, with a median binary neutron star inspiral ranges [2, 43] of 115 Mpc in O3b compared to 108 Mpc in O3a for LIGO Hanford, 133 Mpc in O3b compared to 135 Mpc O3a for LIGO Livingston, and 51 Mpc in O3b compared to 45 Mpc in O3a for Virgo [30]. The duty factor for at least one detector being online was 96.6%; for any two detectors being online at the same time was 85.3%; and for all three detectors together was 51%. Further details regarding instrument performance and data quality for O3b are available in Abbott et al. [30], Acernese et al. [39], Davis et al. [131].

The LIGO and Virgo detectors used a photon recoil-based calibration [83, 207, 357] resulting in a complex-valued, frequency-dependent detector response. Previous studies have documented the systematic error and uncertainty bounds for O3b strain calibration in LIGO [335, 336] and Virgo [38].

Transient noise sources, referred to as glitches, contaminate the data and can affect

the confidence of candidate detections. Times affected by glitches and other data quality issues are identified so that searches for gravitational wave events can exclude (veto) these periods of poor data quality [3, 15, 131, 163, 275]. In addition, several known persistent noise sources are subtracted from the data using information from witness auxiliary sensors [128, 145].

Candidate events, including those reported in Abbott et al. [30] and the new candidates found by the search for sub-threshold counterpart images in Sec. 9.3 of this chapter, have undergone a validation process to evaluate if instrumental artifacts could affect the analysis; this process is described in detail in Sec. 5.5 of Davis et al. [131]. This process can also identify data quality issues that need further mitigation for individual events, such as the subtraction of glitches [115, 129] and non-stationary noise couplings [352], before executing parameter estimation algorithms. See Table XIV of Abbott et al. [30] for the list of events requiring such mitigation.

The GWTC-3 catalog [30] contains 35 events from O3b in addition to the 55 previous events from previous observing runs [33] with a False-Alarm-Rate below two per year, and an expected rate of contamination from detector noise less than 10–15% [30]. We neglect the potential contamination in this analysis. These events were identified by four search pipelines: one minimally modeled transient search cWB [212–214, 216, 217] and three matched-filter searches GstLAL [183, 250, 309], MBTA [40, 50], and PyCBC [42, 43, 123, 276, 351]. Their parameters were estimated through Bayesian inference using the *bilby* [49, 305, 326] and RIFT [221, 289, 371] packages. Both the matched-filter searches and parameter estimation use a variety of binary black holes waveform models which generally combine knowledge from post-Newtonian theory, the effective-one-body formalism, and numerical relativity [for general introductions to these approaches, see 61, 125, 287, 314, and references therein]. The analyses in this paper rely on the same methods, and the specific waveform models and analysis packages used are described in each section.

Of the 35 events from O3b, 31 are likely binary black holes, while four have component masses consistent with being below $3 M_{\odot}$ [24, 30], thus potentially containing a neutron star. We consider these 35 events in the analyses documented in this paper. Specifically, we use the following input data sets for each analysis. The searches for sub-threshold counterpart images in Sec. 9.3 cover the whole O3 strain data set, using the same data quality veto choices as in Abbott et al. [30]

but a strain data set consistent with the parameter estimation analyses: the final calibration version of LIGO data [335] with additional noise subtraction [352]. The posterior-overlap analysis in Sec. 9.3 starts from the posterior samples released with GWTC-3 [182]. The joint-parameter estimation analyses in Sec. 9.3 and microlensing analysis in Sec. 9.4 reanalyze the strain data in short segments around the event times, available from the same data release, with data selection and noise mitigation choices matching those of the parameter estimation analyses in Abbott et al. [30].

9.3 Strong lensing

If a gravitational wave travels close enough to a massive lens, it will produce multiple images, with the number of images depending on the lens profile and source lens geometry. This regime is known as the strong-lensing limit. Each of these lensed images h_j^L will have a change in its amplitude, arrival time and phase compared to the emitted signal h [315]:

$$h_j^L(f) = \sqrt{|\mu_j|} \exp [i2\pi f \Delta t_j - i \text{sign}(f) n_j \pi] h(f), \quad (9.1)$$

for $n_j = 0, 1/2, 1$ for type I, II and III images, which correspond to different minima of the lensing potential. While the magnification μ_j and time delay Δt_j do not affect the waveform morphology (they are completely degenerate with the luminosity distance and coalescence time) the frequency-independent lensing phase shift $n_j \pi$ could induce distortions when the signal has multiple frequency components [119, 154]. In particular, this occurs for type II images since type I does not have a phase shift, and type III only flips the overall sign, which is degenerate with shifting the polarization angle by $\pi/2$. The $\text{sign}(f)$ term is only there to ensure that the time domain waveform is real.

Making a distinction between effects that do and do not change the waveform morphology, we divide our search into two parts. First, we search for pairs of events consistent with the strong-lensing hypothesis. Some of these pairs will have sufficiently strong amplitudes that can be identified as confident detections (super-threshold) by the search pipelines used in Abbott et al. [21, 30, 33], while others may have not been identified as signals (sub-threshold) because of the relative demagnification. Our searches will include both sub- and super-threshold pairs. A pair is the minimum association, but higher multiplicities are also possible. Then, we search for strong lensing focusing on the distortion of type II images.

Sub-threshold Search

In this section, we describe the search for possible sub-threshold counterparts of super-threshold detections from O3. We perform searches over all O3 strain data following the rules for data selection described in Abbott et al. [33] and Abbott et al. [30]. A general search for gravitational waves uses a large template bank covering a broad parameter space as we have no prior information about the signal subspace, resulting in a high trials factor and hence incurring a high noise background. Sub-threshold (lensed) gravitational waves with smaller amplitudes will therefore be easily buried in the noise without being identified as detections as they cannot pass the usual detection threshold.

To uncover these sub-threshold (lensed) signals, we have to effectively reduce the noise background while keeping the targeted foreground (i.e. the signals) constant [122, 223, 247]. The strong lensing hypothesis asserts that lensed gravitational waves, super-threshold or sub-threshold, coming from the same origin have identical waveforms apart from an overall scaling factor and a Morse phase factor as described in Eq. 9.1, and hence should have consistent inferred intrinsic masses and spins.¹ Therefore, we can construct a reduced template bank with only templates that have masses and spins similar to those of a target super-threshold detection. Using the reduced bank lowers the trials factors and noise background and effectively searches for previously unidentified possible sub-threshold lensed counterparts to the target detection. For each known candidate from O3 with a probability of astrophysical origin $p_{\text{astro}} > 0.5$, we create a reduced template bank using their respective public posterior mass and spin samples released with GWTC-3 [182], ensuring that the templates will match well with their respective target events while improving the ranking statistics of the search for similar events, and hence potentially returning new candidates that previously did not reach the threshold $p_{\text{astro}} > 0.5$ in GWTC-3. Details of how the reduced banks are constructed can be found in Li et al. [223].

Given these template banks, we proceed with configurations and procedures as outlined in Abbott et al. [23] to produce a priority list of potential lensed candidates matching each target event, using GstLAL [250, 309] as the search pipeline. The list of candidates obtained is again further vetted using a sky location consistency check detailed in Abbott et al. [23], Wong et al. [366] to ensure the candidates have

¹ The Morse phase factor for different image types has not been considered in the search described here. Should a gravitational wave include detectable higher-order multipole moments, then the Morse phase factor will cause complicated changes to the waveforms, inducing a loss in the search sensitivity.

consistent sky location with the target event. To avoid false dismissal at this step, we only veto candidates with an overlap in 90% credible region of the sky location $O_{90\%CR} = 0$. All candidates with non-vanishing localization overlap are kept for further follow-up with data quality checks as discussed in Sec. 9.2.

In Table 9.3.1, we list the top five candidates from the *individual* targeted searches for counterparts of the detections reported in O3. As in the O3a lensing paper [23], we do not assess in detail the probability of astrophysical origin for each of these. It is also important to note that the reported FARs do not indicate how likely each trigger is a lensed counterpart of the target event, but only how likely noise produces a trigger with a ranking statistic higher or equal to that of the candidate under consideration using these reduced template banks. Similar to Abbott et al. [23], we account for the fact that we have analyzed ~ 332 days of data multiple times for a total of 76 events, and set the False-Alarm-Rate threshold to be 1 in 69 years (i.e. 4.59×10^{-10} Hz). We followed up on the top two candidates listed that passed the False-Alarm-Rate threshold through *goLum*'s joint PE analysis [198, discussed in 9.3]. The results are included in Table 9.3.1. Since both pairs of candidates have mildly *negative* \log_{10} coherence ratios, showing that there is no evidence supporting the lensing hypothesis for either of these pairs, we did not further follow them up with the more computationally intensive *hanabi* analysis [235, discussed in 9.3].

Preliminary Identification of Lensed Pair Candidates

Multiple, non-overlapping images produced by strongly lensed gravitational wave signals have identical phase evolution, and therefore their intrinsic parameters (as well as their orbit's inclination with respect to the line of sight) are expected to have overlapping posteriors. In addition, the angular separations of images (produced by galaxies or galaxy clusters) are several orders of magnitude smaller than the uncertainties associated with their gravitational wave sky location. As a result, their sky localisations will also overlap. (As in the previous section, the Morse phase for different image types is not considered here.)

Under these assumptions, a Bayes-factor statistic (B_U^L) that assesses the consistency between a lensed candidate pair's posterior distributions of intrinsic parameters, sky location, and inclination angle (and thus acts as a discriminator between the lensed and unlensed hypotheses) can be constructed [188]. To convert this statistic to a False-Positive-Probability,² a background distribution of unlensed B_U^L needs to be

² False-Alarm-Rate and False-Positive-Probability, while conceptually similar, pertain to differ-

Table 9.3.1

Top 5 candidates from individual sub-threshold searches for strongly-lensed counterpart images of O3 events from GWTC-3.

Target event	Lensed candidate (UTC)	Δt [days]	FAR [yr^{-1}]	$O_{90\%CR}$ [%]	$\log_{10}(C_U^L)$
GW190930_133541	19-08-05 13:43:48	-56.0	0.002	61.00%	-6.4
GW191204_171526	19-08-05 13:43:48	-121.1	0.006	25.40%	-12.2
GW190828_065509	19-11-12 12:13:18	76.2	0.023	18.00%	-
GW190725_174728	19-08-05 13:43:48	10.8	0.038	39.50%	-
GW190828_065509	20-02-06 07:24:59	162.0	0.154	36.00%	-

The first column lists the target event from O3. The second column shows the time (YY-MM-DD HH-MM-SS) in UTC of the found sub-threshold candidate. The third column shows the time difference (in days) between the candidate and the target event. The fourth column shows the redshifted chirp mass of the template that found the trigger. The fifth column shows the redshifted chirp mass of the target event. The sixth column shows the FARs from the individual search for the new candidate from the second column. The seventh column shows the percentage overlap of the 90% sky localization regions between the candidate and the target event. The eighth column shows the \log_{10} coherence ratio obtained from `golum`'s joint PE analysis.

estimated.

To that end, we conduct an injection campaign involving binary black holes only, in which we sample component masses $m_{1,2}$ from a power-law distribution [4] in the range $(10\text{--}50M_\odot)$. We assume that the redshift distribution of binary black holes is similar to population synthesis simulations of isolated binary evolution [56, 57, 72, 143, 148, 375]. All other parameters are sampled from uninformative prior distributions [188]. We inject the simulated signals into Gaussian noise with O3a representative Power Spectral Density for a LIGO–Virgo detector network. We compute B_U^L for all possible pairs in this injection set, and following Abbott et al. [23], we assign an False-Positive-Probability to a candidate pair using its B_U^L . Candidate lensed pairs involving binary neutron star or neutron-star-black-hole events are not analyzed and ranked.

We additionally employ a Machine Learning-based binary classification scheme to rapidly provide a probability of class membership (lensed or unlensed) for a given candidate binary black holes pair [177]. Such an analysis not only serves

ent contexts in this work. In particular, we use False-Positive-Probability exclusively for significances associated with candidate lensed pairs to discriminate them from unlensed pairs. On the other hand, a False-Alarm-Rate is associated with the significance assigned to individual candidate gravitational wave signal events.

as an independent method to rank candidate pairs but also provides a quantitative significance to pairs for which source-parameter inference samples are unavailable.

Q-transform-based [98] time–frequency maps of strongly lensed binary black holes are expected to have similar shapes, although the signal energy in each time–frequency tile will differ between images. Furthermore, as mentioned earlier, their sky localisations will overlap. Exploiting these facts, Machine Learning models that take Q-transforms and Bayestar [319] sky localisations as inputs, are built. These models use a DenseNet [191] architecture [with several layers pre-trained on the ImageNet dataset; 136], and XGBoost [100] algorithms, trained on lensed and unlensed binary black holes signals injected in Gaussian noise (for details on the Machine Learning training set, see Goyal et al. [177]) The outputs of the individual models are then combined to provide a probability that a candidate pair is lensed or unlensed.

To convert this probability to an False-Positive-Probability, we construct a background distribution of Machine Learning probabilities using a population of unlensed binary black holes events injected in Gaussian noise characterized by the O3a representative Power Spectral Density – the same as was used for the posterior overlap statistic. This Power Spectral Density is found to be sufficiently similar to the averaged O3 Power Spectral Density for the estimation of the background distribution so as not to change the preliminary selection of candidate pairs. The binary black holes population is identical to the one used by the posterior overlap statistic analysis to construct its corresponding background distribution. Furthermore, the sky localisations used to rank candidate pairs come from the same PE analysis used to estimate the posterior overlap statistic ³.

A plot comparing the False-Positive-Probabilities assigned by the posterior overlap and Machine Learning analyses is shown in Fig. 9.3.1. Candidates that have either a posterior-overlap-assigned False-Positive-Probability or Machine Learning-assigned-False-Positive-Probability, (or both), that are smaller than 1%, are selected for more comprehensive Bayesian analyses.

Joint Parameter Estimation

Similar to the analysis of O3a data [23], we perform a joint PE analysis for the most relevant candidate lensing pairs. We follow up on the pairs that display low

³ Note that Bayestar, which is used to assign Machine Learning probabilities to real-event candidate pairs, is expected to provide sky localisations that are similar to those provided by this parameter estimation analysis.




Figure 9.3.1: The False-Positive-Probabilities of each lensed candidate pair constructed from the set of gravitational wave events that exceed an astrophysical probability [159, 206] threshold of 0.5, as evaluated using the B_U^L and Machine Learning classification statistics. Orange dashed lines that correspond to an False-Positive-Probability threshold of 10^{-2} , are also placed. Pairs whose B_U^L -based or Machine Learning-based False-Positive-Probabilities fall below this threshold are selected for additional joint parameter estimation analyses. $B_U^L < 10^{-6}$ has been mapped to an False-Positive-Probability of 1, which is reflected in the gap along the vertical axis between 0.4 and 1.

FPP in their posterior overlap or Machine Learning classification scheme. These are pairs within the whole of O3, but we only consider here those with at least one event in O3b since pairs in O3a were studied in Abbott et al. [23]. We use two complementary pipelines: *golum* [198] and *hanabi* [235]. Both pipelines use the nested sampling algorithm *dynesty* [333], and implement the joint PE with the

help of `bilby` [49, 305].

`golum` [198] is a joint PE tool where the workload is reduced by analyzing the two images, under the lensed hypothesis, in two successive stages. The first image is characterized by the same parameters of the unlensed case (where the time of coalescence and the luminosity distance are the observed ones) with an additional Morse factor. The second image is then analyzed using (samples of) the posterior from the first image as the prior and linking the parameters modified by lensing through three lensing parameters: a time difference, a relative magnification, and a Morse factor difference. The final coherence ratio C_U^L is the ratio of the product of the evidences for the two runs under the lensed hypothesis and the product of evidences for the two images analyzed under the unlensed hypothesis.

`hanabi` [235] first performs a joint inference on a signal pair by constructing a joint likelihood function that is a product of the likelihood function for each individual event, with a joint prior distribution. The latter is defined for a set of joint parameters that can simultaneously describe both signals if they are truly lensed, for example, the masses and the spins, as well as a set of parameters that are different for each of the signals such as the time of arrival, the apparent luminosity distance, and the Morse phase factor associated to each of the lensed signals. The joint parameter space is explored with the package `hanabi.inference` [235]. The inference result is then reweighted with an astrophysically motivated prior distribution; for example, the astrophysical prior distribution for the redshifted component masses would be dependent on both the population model for the intrinsic BBH masses and the redshift distribution of the sources. However, the true source redshift cannot be determined from gravitational wave observations alone since the true source redshift is degenerate with the magnification from strong lensing. To compute the Bayes factor \mathcal{B}_U^L , the source redshift, which serves as a hyper-parameter for the signal pair, must be marginalized over. Selection effects enter as a normalization constant to the marginal data likelihood. This procedure is implemented in `hanabi.hierarchical` with the help of `GWPOPULATION` [338]. The ratio of unnormalized evidences calculated under the lensed hypothesis and the unlensed hypothesis using this astrophysical prior is referred to as the population-weighted coherence ratio $C_{U|pop}^L$, while the ratio of normalized evidences that accounts for both population prior and selection effects is referred to as the Bayes factor \mathcal{B}_U^L in this analysis. We follow our fiducial Singular Isothermal Sphere lensing model when computing the magnification prior [23]. This analysis however does not impose any

informative prior on the time delay or the image types from the lensing model.

Both pipelines use `IMRPHENOMXPHM` [295] as the waveform model, with an additional Morse phase applied to each of the waveform polarizations in the frequency domain. Other inputs, such as the power spectral density estimates and the calibration envelopes, are chosen to match the analyses done in the GWTC-3 catalog paper [30]. Following the same prescriptions of the other analyses, we fix the BBH population model to the Power-Law + Peak model for the primary masses and the merger rate history to Madau–Dickinson star-formation rate [239] normalized by the median GWTC-3 rate [29].

Taking advantage of `golum`’s rapid joint PE, we analyze the 75 pairs of candidates highlighted by posterior overlap and Machine Learning. For each of them, we compute the coherence ratio, which accounts for the probability ratio of the lensed and unlensed hypotheses without including selection effects and population priors. We find that there is a wide range of $\log_{10}(C_{\text{U}}^{\text{L}})$ values, with a peak slightly above zero. This comes from the fact that this analysis considers only triggers already flagged by the posterior overlap and Machine Learning analyses. As a consequence, the analysis is biased towards the higher values. Nevertheless, a significant proportion of events flagged with the Machine Learning pipeline and the posterior overlap pipeline are disfavored, having $\log_{10}(C_{\text{U}}^{\text{L}}) < 0$. When comparing the highest coherence ratio found in the data, $\log_{10}(C_{\text{U}}^{\text{L}}) = 2.5$, with a background of unlensed events, we find that it is well within the expected values, with 1% of the background events having larger C_{U}^{L} . This background is computed for a population of compact binaries that follows the mass, spin and redshift distribution of GWTC-3 [29]. This large number of positive $\log_{10}(C_{\text{U}}^{\text{L}})$ is consistent with the high number of expected false alarms [94, 364]. For those pairs with the highest coherence ratio, we follow up with the `hanabi` pipeline for a total of 17 pairs. Our main results are presented in Fig. 9.3.2, where the left column indicates the event pairs and the horizontal axis their $\mathcal{B}_{\text{U}}^{\text{L}}$. There we can observe that none of the event pairs shows support for the lensing hypothesis, i.e. all $\mathcal{B}_{\text{U}}^{\text{L}} < 1$. The pair with highest $\mathcal{B}_{\text{U}}^{\text{L}}$ is GW190620_030421 – GW200216_220804, for an evidence against lensing of $\sim 1/100$ with the fiducial merger rate density model following the Madau-Dickinson star-formation rate. As a robustness check of how using different merger rate density models would change the results, we repeat the calculations using two more models, namely $R_{\text{min}}(z)$ and $R_{\text{max}}(z)$ from our previous O3a analysis [23] that minimally and maximally bracket many existing population-synthesis results [56, 57, 143, 148]. We see that while the

exact values for the Bayes factor change with the use of different merger rate density models, the conclusion remains that there is no support for the lensing hypothesis in any of the event pairs analyzed. To further assess the significance of these pairs we also include a color code to indicate the probability of having an astrophysical origin $p_{\text{astro}}^{\text{pair}}$, defined as the product of the highest p_{astro} of each event reported in the GWTC-3 catalog paper [30] by different pipelines. In conclusion, we find no evidence of multiply imaged events.




Figure 9.3.2: Bayes factors B_U^L from *hanabi* for the highest-ranked multiple-image candidate pairs. As a check on the robustness of our results, we show the Bayes factors calculated using three different merger rate density models, namely the fiducial model tracking the Madau–Dickinson star-formation rate [239], and also the $R_{\text{min}}(z)$ and $R_{\text{max}}(z)$ model introduced in Abbott et al. [23]. The color for each marker represents the value of $p_{\text{astro}}^{\text{pair}}$ for each pair, which is the probability that both of the signals from a pair are of astrophysical origins and not from terrestrial sources.

Type II image search

In addition to the search for strong-lensing identifying multiple images, we also look for the distortions that lensing introduces in type II images [154]. This is because the frequency-independent phase shift that each image acquires becomes a frequency-dependent time delay for different frequency components. Therefore, for signals containing different measurable spherical harmonic modes, as recently detected in GW190412 [17], GW190814 [19], and other events [30], the overall lensed waveform can be distorted. The extent of the distortion is subject to the power in modes beyond the quadrupole radiation. As a consequence, we do not expect to see these distortions in the majority of the lensed events with current

sensitivities. However, if not searched for, they might be mistaken with deviations from general relativity [156].

To look for these distortions, we use `GOLUM` [198]. Within GWTC-3 we identify 10 events whose posterior has some information about the Morse phase, either by favoring or disfavoring the distortions of the type II image by more than 4% with respect to normality, i.e. the probability of each image type $p(n_j)$ is $p(n_j) > 0.37$ or $p(n_j) < 0.29$. We summarize the evidence of one image type versus another in Fig. 9.3.3. Since only type II images display waveform distortions, we only compute the Bayes factors of the type-II-vs-I and the type-II-vs-III hypotheses. As can be seen in Fig. 9.3.3, only a few events display a preference for one image type versus the other one. This is expected given the signal-to-noise ratio of these events and their power in higher multipole moments. However, GW190412 and GW200129_065458 present higher evidence for type II images. For GW190412 we find a \log_{10} Bayes factor for type II vs. I of 0.60 ± 0.16 and for type II vs. III of 0.22 ± 0.16 . For GW200129_065458 we find 0.38 ± 0.14 and 0.24 ± 0.14 for type II vs. I and type II vs. III respectively. These events have possible super-threshold counterparts but those were discarded by the `GOLUM` analysis. In addition, we have also searched for sub-threshold triggers associated with these events, but found none.

To assess the significance of the type II images, we follow up on GW190412 and GW200129_065458 performing a simulation campaign of type I and type II images. GW190412 simulations show that indeed this event has enough power in higher multipole moments to favor the type II hypothesis so that it could meaningfully test that hypothesis and would favor it if it were true. For GW200129_065458, however, that is not the case. Moreover, GW200129_065458 might have a significant glitch under subtraction [291]. The preference of GW190412 for a type II image could be just a systematic effect due to the waveform modeling, especially since this event falls in challenging parts of the parameter space [17, 109, 185]. For this reason, we repeat the analysis with different waveform families from our fiducial `IMRPHENOMXPHM` model [295]. We find that the preference for a type II image remains when using `SEOBNRv4PHM` [285] or `IMRPHENOMPv3HM` [210]. The same conclusion holds when using different noise realizations for the simulations. Details on these simulation campaigns can be found in Appendix 9.A.

Although we find a mild preference for the type II image hypothesis in GW190412, we find that this analysis cannot provide conclusive evidence of strong lensing. However, our techniques and pipeline will be relevant for future observing runs




Figure 9.3.3: Distribution of Bayes factors comparing different image type hypotheses for the 10 most relevant events. We compare the probability of being type II vs. type I (blue-solid histogram) and of being type II vs. type III (orange-dashed histogram). Only type II images display waveform distortions and for that reason, we do not compare type III vs. type I.

when high-signal-to-noise ratio events display stronger evidence of higher-order modes.

9.4 Microlensing Effects

When the characteristic wavelengths of gravitational waves are comparable to the Schwarzschild radius of a lens ($\lambda_{\text{GW}} \sim R_{\text{Sch}}^{\text{lens}}$), we may observe frequency-dependent magnification of the waveform that can inform us about the lens model [90, 95, 101, 105, 116, 121, 138, 139, 205, 220, 286, 337]. Since the gravitational waves of sources such as binary black holes sweeps through a wide range of frequencies, these beating patterns can reveal the presence of intervening microlenses. In the sensitive range of ground-based detectors, these effects are expected for objects up to $\sim 10^5 M$, which includes stellar-mass objects and intermediate-mass black holes.

Objects that can cause these microlensing effects are predominantly found in larger structures. Therefore we expect that realistic microlensing due to a field of microlenses embedded in an external macromodel potential such as galaxies and galaxy

clusters causes complex effects on the unlensed waveforms [138]. While the effects of these systems on gravitational wave signals have been studied [101, 139, 257, 373], the resulting waveforms are computationally costly to evaluate. Nevertheless, in the absence of specific knowledge of the matter distribution along the travel path and to keep the problem computationally tractable, we assume that the beating patterns are caused by isolated point masses as a first approximation. In this case, the microlensed waveform h^{Micro} can be related to the unlensed waveform h^{U} according to

$$h^{\text{Micro}}(f; \theta, M_{\text{L}}^z, y) = h^{\text{U}}(f; \theta) F(f; M_{\text{L}}^z, y), \quad (9.2)$$

where θ represents the set of parameters defining an unlensed gravitational wave signal, $M_{\text{L}}^z = M_{\text{L}}(1 + z_{\text{l}})$ is the redshifted lens mass, y is the dimensionless impact parameter, and $F(f; M_{\text{L}}^z, y)$ is the frequency-dependent lensing magnification factor [e.g., 337].

Similar to Abbott et al. [23], we perform Bayesian inference on all events from O3b using the unlensed signal model h^{U} and the microlensing signal model h^{Micro} . In particular, we use `bilby` [49, 305] and the nested sampling algorithm `dynesty` [333]. Data products such as strain data and Power Spectral Densities are the same as for GWTC-3 and between the two signal models [30]. For the gravitational wave parameters, we use the same priors as GWTC-3, while the prior on the lens mass M_{L}^z is log uniform in the range $[1-10^5 \text{ M}_{\odot}]$ and the prior on the impact parameter is $p(y) \propto y$ between $[0.1, 3]$. All events were analyzed using `IMRPHENOMXPHM` [295].

The process yields posterior probability distributions of θ or $\{\theta, M_{\text{L}}^z, y\}$ for the unlensed and lensed signal models, respectively. Moreover, we compute the evidence ratio between the microlensed and unlensed signal models, better known as the Bayes factor $\mathcal{B}_{\text{U}}^{\text{Micro}}$.

Fig. 9.4.1 shows the distribution of $\log_{10} \mathcal{B}_{\text{U}}^{\text{Micro}}$ for all the events in O3 and simulated unlensed signals from Abbott et al. [23]. The distribution of $\log_{10} \mathcal{B}_{\text{U}}^{\text{Micro}}$ is primarily clustered around 0 and the distribution for O3 events does not extend to significantly higher values than the distribution for simulated signals. The marginalized posteriors of the microlensing parameters are shown in Appendix 9.B. We conclude that there is no compelling evidence for the presence of microlensing signatures.

9.5 Implications

In this section, we consider some of the implications that derive from the search for lensing signatures. We first forecast the number of detectable strongly lensed




Figure 9.4.1: Distribution of microlensing \log_{10} Bayes factors $\mathcal{B}_U^{\text{Micro}}$ for all events in O3 (blue, solid line) and simulated unlensed signals (orange, dashed line) from Abbott et al. [23].

events based on the latest knowledge on the merger-rate density. Next, we infer upper limits on the strong lensing rate using the non-detection of resolvable strongly lensed binary black holes events. Finally, we use the non-detection of microlensing to infer the compact dark matter fraction in the Universe.

Strong lensing rate

We predict the rate of lensing using the standard methods outlined in the literature [224, 267, 274, 281, 364, 372], at galaxy and galaxy-cluster lens mass scales. To model the lens population, we need to choose a density profile and a mass function. We adopt the Singular Isothermal Sphere density profile for both galaxies and galaxy clusters. Moreover, we use the velocity dispersion function from the Sloan Digital Sky Survey [103] for galaxies and the halo mass function from Tinker et al. [342] for clusters which have also been used in other lensing studies [e.g., 282, 302]. The Singular Isothermal Sphere profile can accurately describe lensing by galaxies, but the mass distribution of clusters tends to be more complicated.

Under the Singular Isothermal Sphere model, we obtain two images with different

magnifications and arrival times. The rate of strong lensing is given by

$$\mathcal{R}_{\text{lens}} = \int \frac{dN(M_h, z_l)}{dM_h} \frac{dD_c}{dz_l} \frac{\mathcal{R}_m(z_m)}{1+z_m} \frac{dV_c}{dz_m} \sigma(M_h, z_l, z_m, \rho, \rho_c) \times p(\rho|z_m) d\rho dz_m dz_l dM_h, \quad (9.3)$$

where $dN(M_h, z_l)/dM_h$ is the differential comoving number density of lensing halos in a halo mass shell at lens redshift z_l , D_c and V_c are the comoving distance and volume, respectively, at a given redshift, $\mathcal{R}_m(z_m)$ is the total comoving merger rate density at redshift z_m , $(1+z_m)$ accounts for the cosmological time dilation, $p(\rho | z_m)$ is the distribution of signal-to-noise ratio at a given redshift, ρ_c is the network signal-to-noise ratio threshold, and σ is the lensing cross-section which indicates, as a function of its various arguments, how efficiently strong lensing will occur. We model the mass distribution of binary black holes following the results for the POWER LAW + PEAK model of Abbott et al. [29]. We consider a merger rate density model that assumes the Madau–Dickinson ansatz [239] that is consistent with recent results from GWTC-3. Moreover, we make use of the absence of a detected Stochastic Gravitational Wave Background to further constrain the merger rate density [29]. For consistency with previous analyses [e.g., 20], we take the Hubble constant from Planck 2015 observations to be $H_0 = 67.9 \text{ km s}^{-1} \text{ Mpc}^{-1}$ [41].

Furthermore, we choose $\rho_c = 8$ as a point estimator of the detectability of gravitational wave signals. We find this choice to be consistent with the search results in Abbott et al. [21] and Sec. 9.3, and we estimate its impact to be subdominant compared to other sources of uncertainty.

Nevertheless, Robertson et al. [302] have demonstrated that the Singular Isothermal Sphere model can reproduce the lensing rate predictions from a study of numerically simulated cluster lenses. Thus, we adopt the same model for both galaxies and galaxy clusters.

In Table 9.5.1, we show our estimates for the relative rate of lensing expected to be observed by the LIGO–Virgo network of detectors. The results are shown separately for galaxy-scale and cluster-scale lenses. Furthermore, these rates are calculated for events that are doubly lensed and for two cases: when only a single event (i.e., the brighter one) is detected (S), and when both of the doubly lensed events are detected (D). The expected fractional rate of lensing (i.e. the lensed to unlensed rate) spans the range $O(10^{-4}–10^{-3})$, depending on the merger rate density assumed. We estimate the fractional rate of observed double (single) events for galaxy-scale

Table 9.5.1: Expected fractional rates of observable lensed double events at current LIGO–Virgo sensitivity.

Merger Rate Density Model	Galaxies		Galaxy Clusters	
	R_D	R_S	R_D	R_S
GWTC-3+Stochastic	$1.9\text{--}11.0 \times 10^{-4}$	$5.0\text{--}19.5 \times 10^{-4}$	$0.8\text{--}4.4 \times 10^{-4}$	$2.0\text{--}7.6 \times 10^{-4}$

This table lists the relative rates of lensed double events expected to be observed by LIGO–Virgo at the current sensitivity where both of the lensed events are detected (R_D) and only one of the lensed events is detected (R_S) above the signal-to-noise ratio threshold. The rates encompass a 90 percent credible interval. We show the rate of lensing by galaxies ($\sigma_{\text{vd}} = 10\text{--}300 \text{ km s}^{-1}$) and galaxy clusters ($\log_{10}(M_{\text{halo}}/M_{\odot}) \sim 14\text{--}16$) separately.

lenses to lie in the range $1.9\text{--}11.0 \times 10^{-4}$ ($5.0\text{--}19.5 \times 10^{-4}$). Similarly, for cluster-scale lenses, the fractional rate is estimated to be in the range of $0.8\text{--}4.4 \times 10^{-4}$ ($2.0\text{--}7.6 \times 10^{-4}$), typically lower than the rates on galaxy scales. These estimates suggest that observing a lensed double image is unlikely at the current sensitivity of the LIGO–Virgo network of detectors. Nevertheless, at design sensitivity and with future upgrades, standard forecasts suggest that the possibility of observing such events might become significant [224, 267, 274, 281, 364, 372]. Compared with other lens models, our lensing rates are consistent with those predicted for Singular Isothermal Ellipsoid models [e.g., 281, 364, 372].

Implications from the non-observation of strongly lensed events

The absence of any detections of strongly lensed gravitational wave events before and during O3 provides a complementary way to constrain the merger rates of compact objects at high redshift. The detection of individual gravitational wave events has enabled measurement of the low redshift ($z < 1$) merger rate [29]. However, the high redshift merger rate of gravitational wave sources is not yet measured directly, and we have only been able to place an upper limit on it from the absence of a detection of the Stochastic Gravitational Wave Background [20]. The absence of such a detection naturally leads to a bound on the lensing rate expected from GWTC-3 [81, 266].

By using the same power-law form for the merger rate as that used in Sec. 9.5, but extended up to $z = 2$, we obtain limits on the merger rate at redshift $z > 1$ from the absence of detections of strongly lensed events. The corresponding constraints (90% credible intervals) are shown in Fig. 8.3.1 as the cross-hatched region bounded

by the dash-dotted curves. The changes in the upper bound of the merger rates are driven by the absence of detected lensing events, whereas the lower bound is driven by the low-redshift constraints on the merger rate. For comparison, the current limits on the merger rate from GWTC-3 up to redshift $z = 1$ [29], with the bounding curves extrapolated to higher redshifts $z > 1$, are shown as the grey shaded region bounded by the dotted curves. For further comparison, we have also plotted the solid black curves which show the current constraints from the absence of detection of the Stochastic Gravitational Wave Background [20]. The upper bounds on the merger rate from lensing are more stringent than the bounds from GWTC-3 at high redshift [29], and are also comparable with the bounds from the Stochastic Gravitational Wave Background for redshifts $z < 1.2$. The slight difference between the constraints on the merger rates at low redshift derived from the Stochastic Gravitational Wave Background [20] and from GWTC-3 [29] arise because the bounds from the Stochastic Gravitational Wave Background are obtained here using the previous constraints on the merger rate at low redshift derived using GWTC-2 [22].

Constraints on compact dark matter from gravitational-wave microlensing

Black holes and other compact objects have sizes comparable to their gravitational radius, and may cause microlensing effects on GW signals. Although their abundance is heavily constrained by several astronomical observations [91, 92], the possibility of their contributing to dark matter cannot be ruled out in several mass windows.

Here we use the non-observation of microlensing effects on the gravitational wave signals detected by LIGO and Virgo to constrain the fraction of dark matter contributed by compact objects in the mass range $\sim 10^2\text{--}10^5 M_\odot$ [53, 205, 350]. The essential idea is that if a significant fraction of dark matter is in the form of compact objects, they would introduce detectable microlensing signatures on the gravitational wave signals that we observe.

Assuming that lensed and unlensed events occur as Poisson processes, we compute the posterior distribution on the lensing fraction ($u \equiv \Lambda_\ell/\Lambda$), defined as the ratio of Poisson means of lensed events to the total number of detected events. This is then used to compute the posterior of the fraction of compact dark matter ($f_{\text{DM}} \equiv \Omega_{\text{CO}}/\Omega_{\text{DM}}$) [53]. We take that a total of $N = 67$ binary black holes mergers are




Figure 9.5.1: Merger rate density as a function of redshift based on the GWTC-3 results without lensing constraints (grey) and with lensing constraints (cross-hatching) included. For clarity, we show only the results for galaxy-scale lenses. Because lensed detections may occur at higher redshifts than unlensed events, their non-observation can be used to constrain the rate of mergers at higher redshifts. The ‘No lensing’ results shown here do not include constraints derived from the absence of an Stochastic Gravitational Wave Background. The latter constraints are shown separately by the solid black curves.

detected during the O3 run ⁴, and none of them is lensed (i.e., $N_\ell = 0$). We then estimate the posterior distribution of the lensing fraction u . Finally, the posterior of f_{DM} can be computed as

$$p(f_{\text{DM}} | \{N_\ell = 0, N\}) = p(u | \{N_\ell = 0, N\}) \left| \frac{du}{df_{\text{DM}}} \right|, \quad (9.4)$$

where du/df_{DM} is the Jacobian that relates the observed fraction u of lensed events to the compact dark matter fraction f_{DM} in the Universe.

We determine this Jacobian by simulating astrophysical populations of binary black holes mergers lensed by point mass lenses [53]. ⁵ The constraints we obtain depend upon the assumed distributions of the component masses, spins and the redshifts of the mergers, which have considerable uncertainties. We assume that the masses are

⁴ These are the events cataloged in GWTC-3 that do not contain a neutron star component.

⁵ The simulations are done assuming the O3b representative Power Spectral Density and Gaussian noise. The Jacobian is not expected to change significantly if real noise is used instead.

distributed according to the POWER-LAW + PEAK model of Abbott et al. [29] while spins are assumed to be aligned/antialigned with the orbital angular momentum with magnitudes distributed uniformly in $(0, 0.99)$. We use the approximant IMRPHENOMD [209] to produce the waveforms. We consider different redshift distributions of the mergers: uniform distribution in comoving volume, the power-law model of Abbott et al. [29], the Madau-Dickinson model [239], as well as some representative population-synthesis models given by Dominik et al. [143] and Belczynski et al. [58]. In our simulations, compact objects are approximated by point mass lenses and distributed uniformly in comoving volume. Binaries producing a network signal-to-noise ratio of 8 or above in the LIGO–Virgo detectors are deemed detectable. In order to reduce the computational cost of performing the simulations, we estimate $\mathcal{B}_U^{\text{Micro}}$ using an approximation to the Bayes factor that is expected to be accurate in the high-signal-to-noise ratio regime [113, 353]. We then compute the fraction of detected events that produce a $\mathcal{B}_U^{\text{Micro}}$ larger than the highest $\mathcal{B}_U^{\text{Micro}}$ obtained from real LIGO–Virgo events. This lensing fraction is computed as a function of the f_{DM} , which is used to compute the Jacobian du/df_{DM} .

The largest value of the microlensing likelihood ratio obtained from GWTC-3 events is $\log_{10} \mathcal{B}_U^{\text{Micro}} = 0.799$. We compute the fraction of simulated events with $\log_{10} \mathcal{B}_U^{\text{Micro}} \geq 0.799$, for different lens masses. This allows us to compute the Jacobian du/df_{DM} and thus the posterior on f_{DM} . The 90% upper limits are shown as a function of the lens mass (assuming a monochromatic spectrum) in Fig. 9.5.2. A number of bounds on f_{DM} , for masses comparable to the mass range considered in this analysis, have been estimated, notably from EM searches for microlensed events. These include constraints from searches for lensed FRBs [112, 226, 262], GRBs [192, 229, 292], radio emissions from active galaxies [355]; from searches in archival EROS-2 and MACHO photometric data for Galactic microlenses [60]; and from searches for microlensed quasars [150] and supernovae [169, 377]. The bounds we obtain are weaker than several of these constraints, some of which are well below 1%, as compared to the $O(10\%)$ constraints obtained from our analysis. Nevertheless, GWs are a fundamentally different messenger than EM waves, and do not suffer from the same uncertainties and systematics. Thus, acquiring constraints from the non-detection of GW microlensing is a worthwhile exercise, even though presently they are modest.

The weak constraints can be attributed to an insufficient number of BBH detections. Indeed, the gravitational wave lensing bounds will improve significantly in the next




Figure 9.5.2: The spread in the 90% upper limits on f_{DM} obtained from the O3 events using 5 different redshift distribution models for binary black holes mergers: Belczynski et al. [58], Dominik et al. [143], Madau and Dickinson [239], Abbott et al. [29] and uniform in comoving 4-volume, assuming a monochromatic mass spectrum for the compact objects forming dark matter. The lens mass is shown on the horizontal axis. The grey (black) shaded regions correspond to the spread in f_{DM} upper bounds computed assuming flat (Jeffreys) prior on Λ and Λ_ℓ . The upper and lower curves bounding the spreads correspond to the most pessimistic (weakest) and optimistic (strongest) upper limits, as determined from the set of assumed redshift distributions, in each mass bin. The current f_{DM} constraints are weaker relative to other corresponding EM constraints. We refer the reader to [91, 92, 112] for comparison. Nevertheless, these constraints are expected to improve significantly with the increased detection of unlensed BBHs in forthcoming LVK observing runs.

few years as the sensitivity of gravitational wave detectors improve [9] resulting in a sizeable increase in observed BBHs. Assuming ~ 300 binary black holes detections in O4 and $O(1000)$ detections in O5, the constraints on f_{DM} will improve to $\sim 10^{-1}$ and $\sim 10^{-2}$, respectively.

9.6 Concluding Remarks

We have extended the search for lensing signatures to all binary black holes candidates with a probability of astrophysical origin higher than 0.5 from O3b [30]. While we have not observed any significant candidates for strongly lensed events, we updated the constraints on the rate of such events from several different analyses. First, we searched for sub-threshold repeated signals associated with super-threshold events using reduced template banks produced from the posterior probability distributions of the super-threshold events. Interesting sub-threshold/super-threshold pairs and pairs formed from two super-threshold events were further analyzed for their probability of being from a single, strongly lensed source. For super-threshold/super-threshold pairs, we calculated the degree of overlap between the posteriors of the intrinsic parameters and sky location, which were obtained from Bayesian inference. Moreover, we analyzed these pairs using a new analysis based on the comparison of spectrograms through machine learning. Finally, pairs with false-positive probability from either analysis smaller than 10^{-2} were further studied by conducting full joint Bayesian inference analyses that take population priors and selection effects into account. We found no pairs that show significant evidence for strong lensing.

The events from O3b were also analyzed for distortions caused by the lens on the gravitational waveform. First, we searched for the distortions that lensing introduces on type II signals, which are in the form of a frequency-independent phase shift (Morse phase). The Bayes factors for all events show no evidence for type II signal distortions. Similarly, we searched for the frequency-dependent distortions caused by point masses. None of the computed Bayes factors show any significant signs of microlensing. For both analyses, some events show interesting features in the posteriors for the Morse phase or lens mass. However, follow-up analyses using simulated signals show no further signs of the lensing nature of these features. Altogether, we found no significant evidence for distortions of the gravitational waveforms that can be attributed to lensing.

The lack of evidence for lensing is then used to infer properties of the lensing rates and to set constraints on the dark matter fraction of (dark) compact objects.

Finally, we note that our conclusions are based on estimates and assumptions that are in line with other analyses from the LIGO–Virgo–KAGRA Collaboration [29, 30]. It is possible to arrive at different conclusions and interpretations if assumptions are chosen differently. Examples include claims that almost all detections are strongly

lensed if one assumes that heavy black holes do not exist [74, 75, 77]. Data from the upcoming observing runs are expected to further expand the catalog of gravitational wave detections that can further shed light on the lensing of gravitational waves [9]. Moreover, multi-messenger astronomy may provide significant input in confirming and interpreting possible lensed gravitational wave signals [363].

APPENDIX

9.A Type II simulation campaigns

Given the mild evidence of GW190412 and GW200129_065458 towards being a type II strongly lensed image presented in Sec. 9.3, we follow up on these events by doing an injection campaign where we simulate type I and type II images similar to the events, and verify whether the posteriors recovered are compatible with the distribution observed for the real events. These injections are performed in different noise realizations and with different waveform models. The observed feature could be caused by two main other effects than a type II image: noise artifacts or systematic effects in the waveform modeling. In the former, non-Gaussianities in the noise could be such that they lead to the observation of spurious features, while in the latter case, the specific combination of observed parameters could lead to some systematic issues in fitting with the waveform model. Waveform systematics might be especially important for these events since they lie in challenging parts of the parameter space [17, 109, 185]. Moreover, for GW200129_065458 Payne et al. [291] reports that there could be a significant glitch under subtraction.

To test for the noise-related features, we generate colored Gaussian noise from the PSD around the time of the candidate and then inject the maximum likelihood parameters coming from the parameter estimation and take the Morse factor to be either the value for a type I or a type II image. In the first step, we do this for only one noise realization for each event to see whether we can reproduce similar features or not. For GW200129_065458, the injection shows that the effect is too weak to be distinguishable from one image type to the other, as can be seen in the uninformative posteriors of Fig. 9.A.1. As a consequence, no further investigation is done into this event. On the other hand, for GW190412, the feature seen in the real data is compatible with the one seen in the injection (see Fig. 9.A.2).

Given that the real-data results are compatible with the type II injection for GW190412, we investigate further the noise hypothesis. For this purpose, we take the maximum likelihood parameters and a Morse factor of 0 or 1/2 and inject the signal generated with the IMRPHENOMXPHM [295] model in ten different noise realizations. We then repeat the analysis in the same way as for the real signal and verify if we retrieve the same preference for a type II image. For all the noise realizations used here, we see the same behavior as in Fig. 9.A.2.




Figure 9.A.1: Posterior distribution of the Morse phase for GW200129_065458. We compare the real event posterior (solid-blue) with an injection campaign of type I (dashed-orange) and type II (dotted-green) images. Type II images correspond to $n_1 = 1/2$. For this event, the differences between the distribution are small and make it difficult to learn anything additional about the event. The Kolmogorov–Smirnov statistic is 0.07 for type I vs real, and 0.08 for type II vs real.

We perform an extra test by injecting the maximum likelihood parameters with a given image type in the generated noise for different waveform models. We use the IMRPHENOMPv3HM [210] and the SEOBNRv4PHM [285] model to generate the signal and use the IMRPHENOMXPHM [295] model to recover it. This enables us to combine the two possible sources of systematics. This way, we can verify whether a different noise combined with a different model also leads to a preference for type II images. For all the noise realizations and the two models used for the injections, we find that the injections always recover the correct hypothesis, and the fact that the real event supports type II is unlikely to be a result of noise or waveform artifacts, as shown in Fig. 9.A.3.

Although these tests do not discard the type II image hypothesis, they cannot conclusively confirm it. To confirm the presence of lensing for this event with a mild preference for a type II image, we would need additional evidence. Therefore, we search for possible sub-threshold counterparts with the methodology explained in Sec. 9.3. However, we find only marginal triggers.




Figure 9.A.2: Posterior distribution of the Morse phase for GW190412. We compare the real event posterior (solid-blue) with an injection campaign of type I (dashed-orange) and type II (dotted-green) images. Type II images correspond to $n_1 = 1/2$. For this event, the peak seen in the real data and the one seen for the type II image are compatible, hinting at a possible type II image. In this case, the Kolmogorov–Smirnov statistic is 0.20 for type I vs real, and 0.13 for type II vs real.

In the end, these additional searches did not enable us to find any extra evidence for lensing, while still not ruling out the possibility for GW190412 to be a type II image.

9.B Marginalized Posteriors of Microlensing Parameters

As a supplement to the distribution of \log_{10} Bayes factors $\mathcal{B}_U^{\text{Micro}}$ shown in Fig. 9.4.1, we show the individual marginalized posterior distributions of redshifted lens mass M_L^z and $\log_{10} \mathcal{B}_U^{\text{Micro}}$ (right vertical axis) in Fig. 9.B.1. The Bayes factors individually do not show clear evidence for microlensing by point-mass lenses. However, several events show a narrow posterior distribution of the redshifted lens mass. An example is GW200208_130117 (with $\log_{10} \mathcal{B}_U^{\text{Micro}} = 0.8$), for which the waveform corresponding to the maximum posterior for this event, with and without lensing, is shown in Fig. 9.B.2. The beating pattern introduced by the point-mass lens is most visible as a reduction of the amplitude for two cycles in the middle of the signal and an increase in the amplitude before and after this reduction. We hypothesize




Figure 9.A.3: Comparison of the Morse factor distribution for the real event (solid-purple) with the recovered posterior distribution for an injection made with IMRPhenomPv3HM for a type I image (dashed-blue) and a type II image (dotted-orange), and with an injection made with SEOBNRv4PHM for a type I image (dashed-green) and a type II image (dotted-red). In all the cases, the posterior distributions agree with the injected data, with the real event resembling a type II image.

that short-duration noise fluctuations may have caused an apparent dip in the signal, which in turn may have led to a distortion similar to a point-mass lens beating pattern. This is corroborated by a low Bayes factor $\mathcal{B}_U^{\text{Micro}}$, which concludes the data is inconclusive about the microlensing hypothesis.




Figure 9.B.1: Marginalized posterior distributions of redshifted lens mass M_L^z and $\log_{10} \mathcal{B}_U^{\text{Micro}}$ between microlensed and unlensed hypotheses.




Figure 9.B.2: The time-domain waveform corresponding to the maximum posterior of GW200208_130117, with and without the microlensing hypothesis.

FOLLOW-UP ANALYSES TO THE O3 LIGO-VIRGO-KAGRA LENSING SEARCHES

Note: This Chapter is an adaption of the publication:

J. Janquart, ...(6 authors)..., **Alvin K. Y. Li**, ...(22 authors)..., “Follow-up Analyses to the O3 LIGO-Virgo-KAGRA Lensing Searches”, *Mon. Not. Roy. Astron. Soc.* 526.3 (2023), doi:10.1093/mnras/stad2909.

Alvin K. Y. Li led the analysis and provided results for the search for GstLAL-based sub-threshold lensed gravitational waves described in the paper. The other related sections in this chapter are contributions from co-authors.

10.1 Introduction

The LVK collaboration has searched for strong lensing and for microlensing signatures in the following LVK observing runs: O1–O2 [186], O3a [23], and the full O3 run [32], yielding no confident signatures. In parallel, other searches have been performed, confirming that no lensing features have been confidently detected so far [122, 223, 233, 247].

Nevertheless, in these searches, interesting candidates have been found. This is the case, for example, for GW190412 that shows some support for being a type II image, the GW191103–GW191105 pair for strong lensing—discarded only after the inclusion of both the population priors and selection effects—, and GW200208_130117 which displays some features which are compatible, at low significance, with microlensing [32]. Although, ultimately, not confirmed as lensed, such events contain features representative of signatures one could find in genuinely lensed events. It is therefore important to see what sort of follow-up analyses one could do on such events to have a better grasp on their significance, and to extract a maximum of information about the systems.

Possible avenues to achieve this in the future are presented in this work by applying them to these interesting O3 candidates. We follow up on strongly-lensed candidates by making a background distribution of simulated unlensed events in order to compute each candidate’s false-alarm probability (FAP). We also compare

the candidates to the most recent simulations as to see if we can identify the lens that could be at the root of such a lensed event. Additionally, we look for lensed electromagnetic (EM) counterparts by cross-matching with galaxy-lens catalogues. Moreover, since some fainter counterparts are likely to be present in a strongly-lensed multiplet, we also follow up on an additional strongly-lensed candidate containing a supra-threshold event GW191230_180458 and a weaker “sub-threshold” event LGW200104_184028 identified for investigation by a new method [180]. We analyze this pair in more details in this work, showing that it is an intriguing pair but is unlikely to be lensed. We also analyse the most significant candidate microlensing events using different lens models, inferring the parameters of the lens models at the same time. We compare these models to investigate which is most likely. Moreover, we analyze the most significant of these microlensing candidates with a millilensing framework to see if the signatures could come from a source in this lensing regime. We do not report any additional evidence for lensing but outline some important next steps to further deal with a possibly lensed event.

We stress that whilst the events discussed in this paper may be treated as though they were lensed, they do not display significant evidence for lensing [32]. The goal of this work is to demonstrate the methodologies that can be used to dig deeper in the case of genuinely lensed events and to better assess the importance of candidates. To represent this, we refer to the events as “lensed candidates” in what follows. Additionally, since the events and event pairs analyzed in this work have been selected because they present interesting features, it is often the case that they lead to higher Bayes factors. However, this is generally not enough to claim lensing, and we would also require to have posteriors converging to a given value of the lensing parameters or a high significance compared to a background before considering an event as lensed.

The rest of the Chapter is structured as follows. In Sec. 10.2, we introduce the methodologies that are applied in the analysis of this work. In Sec. 10.3, we show the results for different new analyses performed on GW190412, an event flagged with some support for a type II image. Next, in Sec. 10.4, we analyze the GW191103–GW191105 event pair, found to have some characteristics resembling the ones expected for strongly-lensed event pairs. We continue in Sec. 10.5 by analyzing another event pair, GW191230–LGW200104, a pair made of a supra and a sub-threshold event. In Sec. 10.6, we analyze GW200208, an event found to have a mild support for microlensing in past searches. We then give our conclusions and

prospects in Sec. 10.7.

10.2 Analysis Methods

Strong Lensing

When the GW travel path is close enough to a massive lens, gravitational lensing leads to several—possibly detectable—images having the same frequency evolution. This is called the strong lensing regime. In this work, several methods are used to analyze the interesting lensing candidates and compute Bayes factors for lensed versus unlensed hypotheses: $\mathcal{B}_U^L = \mathcal{Z}_L/\mathcal{Z}_U$, where $\mathcal{Z}_H = P(d_1, d_2|H)$ is the evidence under the hypothesis H ($H = L$ for lensed or $H = U$ for unlensed), and $d_i = n_i + h_i^H$ is the data stream for the i^{th} image, made of a noise (n_i) and a GW (h_i^H) component. In this work, we adopt the same conventions as those used in Janquart et al. [198], Lo and Magana Hernandez [235], referring to the evidence ratio as the Bayes factor when it includes both the population and selection effects, and the coherence ratio when these are not included.

The first analysis method is called posterior overlap [PO, 188]. Since the frequency evolution of lensed images is unchanged, the detector frame parameters should match (except for those modified by lensing), there should be a significant overlap between the posteriors obtained for these images under the unlensed hypothesis. Therefore, one can compute a detection statistic comparing the evidence for the lensed and unlensed hypotheses

$$\mathcal{B}^{\text{overlap}} = \int d\theta \frac{P(\theta|d_1)P(\theta|d_2)}{P(\theta)}, \quad (10.1)$$

where $P(\theta|d_i)$ is the posterior obtained from data i , and $P(\theta)$ is the prior.

Typically, kernel-density estimators (KDEs) are performed on the posteriors in Eq. (10.1) for a subset of parameters (component masses, spins, inclination angle and sky location), and $\mathcal{B}^{\text{overlap}}$ is computed using those KDEs. The posteriors used often come from the usual unlensed PE [49, 356].

Another method used is joint parameter estimation (JPE), where one performs the joint inference of the lensed images, linking them through the lensing parameters [122, 198, 202, 233, 235, 247]. These pipelines have two different ways to tackle the problem. Some compute the full joint evidence $p(d_1, d_2|H_L)$ at once, such as those outlined in [233] and [235]. Here, we use the HANABI pipeline from [235]. The alternative approach is to instead consider the evidence for the second image as conditional on that of the first [198, 202]. This makes the computation faster and is

equivalent to JPE under the lensed hypothesis. However, some level of approximation is added by doing sub-sampling to improve the speed. The pipeline undertaking this method used within this work is called GOLUM [198, 202]. The analysis done using the pipeline for the first image also directly offers the possibility to search for type II images when higher-order modes are present [32, 154, 199, 235, 358, 360].

In addition to matching purely on the observational parameters of the system, one can also use models of the lens to inform the strong-lensing detection process [188, 201, 235, 236, 259, 364]. Lensed events do not only have matching frequency-domain evolution but they are also linked via the lensing parameters. Their measured values can be used to assess how likely it is for the observed events to be lensed for a given model. To do this, one compares the probability of having the apparent lensing parameters under the lensed and unlensed hypotheses. This may be done for all of the lensing parameters, or a subset of them. To obtain the probabilities, an unlensed population of BBH mergers is constructed using given population models (mass, redshift, spin, . . . distributions) and the phase differences, relative magnifications, and time delays are computed between these events. In parallel, the same process is performed on a lensed population, produced from a BBH population and a lens population following a specific model such as a Singular Isothermal Sphere (SIS [365]) [188] or a Singular Isothermal Ellipsoid (SIE [218]) [259, 364] for the galaxy lenses. From the two computed distributions, a probability density function can be obtained via, for example, KDE reconstruction. It is then possible to evaluate the ratio

$$S^{\text{gal}} = \frac{P(\Phi|H_L)}{P(\Phi|H_U)}, \quad (10.2)$$

where H_L and H_U designate the lensed and unlensed hypotheses respectively, and Φ is the set of lensing parameters under consideration.

Specific examples of the statistics computed with this method for this work are \mathcal{R}^{gal} [188] using specifically the time delay, and \mathcal{M}^{gal} [259] which may use all or a subset of the lensing parameters. Both models can be used either with an SIS or an SIE lens model. These statistics are used to select candidates to be followed up by the more extensive analyses or are multiplied with the detection statistics evaluating the match between the parameters. Though model dependent, this in general decreases the risk of false alarm detections [188, 201, 364].

Alternatively, one can consistently incorporate the information from a lens and a source population model [235], where the lens and the source population model affect both the probability of observing a given set of data, in this case (d_1, d_2) ,

under the lensed and the unlensed hypothesis. Specifically, the lens population model informs the joint probability distribution on the magnification, the image type, and the time delay between images, as well as the optical depth for strong lensing, while the source population model informs the distribution of the (true) redshift and the source parameters of a lensed source. This was already done in [23] and [32] using the simple SIS lens model.

In practice, it is difficult to write down an analytical form for the above-mentioned joint probability distribution from a lens model except for some simple lens models (e.g. the SIS model), and instead one usually resorts to constructing a surrogate that approximates the probability density function, such as the aforementioned KDE technique. However, it can be computationally expensive to use KDE-based schemes to construct an estimate for the probability density from a catalog of simulated lensed images that contains many (e.g. millions of) samples, which in turn is evaluated over a set of (roughly tens of thousands of) posterior samples.

In this work, we use the probability density surrogate described in [236] that fits the joint probability density on the magnification and the image type conditioned on the time delay between images from a catalog of mock lens images used in [281] using a normalization-flow-based method [234]. The underlying strong lensing model adopted in the simulation is a population of galaxy-scale SIE lenses with external shear. The lens-redshift-dependent velocity dispersion function is constructed from hybridizing the velocity dispersion measurement for the local Universe derived from the Sloan Digital Sky Survey Data Release 6 [59] with the Illustris simulation result for the velocity dispersion function at higher lens-redshifts [344]. The ellipticity and the external shear follow a Gaussian distribution and a log-normal distribution respectively with additional detail found in [281].

Millilensing

When the mass of the lens and its extent is reduced, the different images produced by lensing can start overlapping. In such a case, they interfere, and one gets only one image exhibiting beating patterns. This is called millilensing [230],

To search for such signals, one needs to analyze the GW signal assuming that several lensed images are interfering with each other. Usually, the number of signals is not known beforehand. Therefore, it can either be fixed in the search or it can be a variable one tries to infer [230].

Microlensing

For lens-source systems such that the wavelength of the GW is comparable to the Schwarzschild radius of the lens, frequency-dependent modulation of the amplitude occurs.

In Abbott et al. [23, 32], Hannuksela et al. [186] the microlensing search has been performed using an isolated point mass model. Regardless of the exact model of the lens, the lensed (h_{Micro}) and unlensed (h_{U}) waveforms are linked as

$$h_{\text{Micro}}(f; \theta, M_{\text{Lens}}^z, y) = h_{\text{U}}(f; \theta) \times F(f; M_{\text{Lens}}^z, y), \quad (10.3)$$

where, θ are the parameters defining the unlensed GW signal, y is the dimensionless source position, M_{Lens}^z is the redshifted lens mass, and $F(f; M_{\text{Lens}}^z, y)$ is the amplification factor which is dependent upon both the frequency and the mass-density profile of the lensing object (more details can be found in [337] for example), hence on the lens model.

Whilst the isolated point mass model has been used for its simplicity and consequent speed, there are other density profiles that may describe microlenses. These include, but are not limited to, the SIS [365], the SIE [218], or the Navarro-Frenk-White (NFW) [272] profiles. Efforts have been made to incorporate some of these models into microlensing GW analyses [370] to determine more information about the lensing object in the event of microlensing detection. This work will use these models to analyse the data and recover information about potential lenses that could be at the origin of a lensed event system.

To explore these multiple models, microlensing candidates are analysed using GRAVELAMPS [370]. This algorithm performs PE analyses of the GW data by jointly inferring the source and lens parameters, assuming a particular lens model. Therefore, it can not only extract information on the lens, but also perform lens-model selection by comparing the evidence obtained for different models and see which one is favoured based on the data. In addition, to cross-check the results obtained with this pipeline, the data is also analysed with the GWMAT pipeline [255], a similar but independent PYTHON/CYTHON analysis package implementing the point lens model.

10.3 GW190412

GW190412 [17] is a well-known event as it is, alongside GW190814 [19], one of the events with significant higher order modes (HOMs) detected during O3. It is

identified as the coalescence of a $\sim 30 M_\odot$ black hole, with a $\sim 8 M_\odot$ one. During the O3 lensing search [32], this event was found to be the most likely candidate for being a type II image. However, the evidence ($\log_{10}(\mathcal{B}_I^{\text{II}}) = 0.61$ and $\log_{10}(\mathcal{B}_{\text{III}}^{\text{II}}) = 0.30$) is inconclusive and could be related to other effects as well, such as noise or waveform effects. In this section, we investigate possible origins of this feature. In particular, we re-analyze the data searching for type II images using other waveform models, and see if the observed feature could be related to microlensing effects.

Check for Waveform Systematics

For any astrophysical inference about CBCs from GW data, it is crucial to understand the possible systematic errors due to approximations in the waveform models used or whether observed features could not originate from the model used. Since full numerical relativity simulations are only available as a reference for a few points in parameter space, the most convenient way to study the impact of waveform systematics is to compare results from different models. PE runs for GWTC data releases have always used at least two waveforms from independent modelling approaches and additional studies on events of special interest regularly compare larger numbers of models [see, e.g., 13, 18, 109, 152, 185, 246].

In Abbott et al. [32], the type II image (Morse factor of 0.5) search was performed with the IMRPhenomXPHM waveform [295], including HOM and precession effects. A feature similar to the observed one —meaning a peak around a value of 0.5 in the Morse factor posterior— was recovered in various scenarios. For example, by injecting a signal with the maximum likelihood parameters in simulated noise with a given waveform family and recovering with IMRPhenomXPHM. This seemed to indicate that the feature was consistent with a real signal, at least given the used waveform.

Here, we re-analyse the data using two other commonly-used waveforms: IMRPhenomPv3HM [210] and SEOBNRv4PHM [285]. These two waveforms encapsulate the same type of physics as IMRPhenomXPHM with HOMs and precession included. The analyses are performed using the same setup as the one used for the Type II image search performed in Abbott et al. [32].

The posterior recovery of the Morse factor (n_1), allowing it to take any value from 0 to 1 instead of discrete, for the different waveforms is shown in Fig. 10.3.1. The posterior peaks at 0.5 using IMRPhenomXPHM. When analyzing the data with the IMRPhenomPv3HM waveform we still observe a peak but it is less prominent. On

the other hand, with the SEOBNRv4PHM waveform the posterior distribution has an earlier maximal value, is much broader, and is lacking a pronounced peak. This is different from the results seen with the Phenom-family waveforms. Therefore, the observed feature seems to come from a combination of noise and waveform effects, and the event is probably not a type II image.




Figure 10.3.1: Recovery of the Morse factor for the GW190412 event with different waveforms: IMRPhenomXPHM, IMRPhenomPv3HM, and SEOBNRv4PHM. The support for a type II image is present for the two waveforms from the Phenom family. However, the feature is less prominent for the IMRPhenomPv3HM waveform, and only marginally present for the SEOBNRv4PHM waveform. Therefore, the observed feature is probably spurious and the event is not a type II image.

Microlensing Analysis

The possible signs of de-phasing that generated initial interest in GW190412 as a type II lensing candidate may also be a mistaken interpretation of frequency-dependent microlensing effects. This motivates an analysis of the event using the GRAVELAMPS pipeline to determine if there is any potential favouring of this interpretation of the signal's features.

Looking at the results of the analysis for GW190412 shown in Fig. 10.3.2 ¹ marks

¹ The redshifted lens mass $M_{\text{Lens}}^z (1 + z_{\text{Lens}}) M_{\text{Lens}}$ where M_{Lens} is the lens frame lens mass. In the GRAVELAMPS analyses the redshift of the lens is calculated based on the lens being half way between

the event as unique amongst those analysed for this paper in that it favours the point mass lensing model over the SIS model with $\log_{10}(\mathcal{B}_U^L)$ values of 0.6 and 0.4 respectively. This remains quite marginal preference for the microlensing hypothesis in either case, although it is worth noting that this support is higher than for GW191103 or GW191105, the events analysed in a strong-lensing context in Sec. 10.4. This is consistent with the apparent signs of de-phasing being present only in GW190412. Whether the correlation holds would need to be confirmed with additional examinations of simulations or additional type II lensing candidates.

However, whilst support for this event is higher in terms of the raw Bayes factors, the posteriors for the lensing parameters need to be examined. Fig. 10.3.2 presents these posteriors for the marginally more preferred point mass lensing model. The source posteriors are tightly constrained but the lensing parameters are extremely broad, leading to the conclusion that this event does not indicate any particular signs of microlensing either, and again the apparent features could be related to noise or waveform issues.

10.4 GW191103–GW191105

GW191103 and GW191105 were BBHs detected during the third observing run [30]. In the main LVK analyses, the standard treatment of the signals revealed nothing out of the ordinary for these events. However, when treating the events as potential lensing candidates, the pair display some intriguing characteristics. There is a notable amount of overlap between some of the reported source parameters, such as the sky location and masses; see Fig. 10.4.1 for a representation of the posteriors. Moreover, the two events have about two days delay between their merger times which is consistent with galaxy-scale lenses [259, 364]. However, in the LVK lensing search, these events were ultimately discarded once the JPE Bayes factor had been computed [32], meaning that the observed overlap is unlikely to be coming from a lensed BBH and is more likely to be coincidental. Nevertheless, as was stated in the introduction, in the following analyses we have disregarded this and treated the event as though it were a lensed pair.

Posterior Overlap Candidate Identification

During the O3 strong lensing search, the PO analysis and a machine learning pipeline, LENSID [177], were used to identify potential lensing candidate pairs from the detected events. The top 1% of these—approximately a hundred pairs—were

the source and the observer.

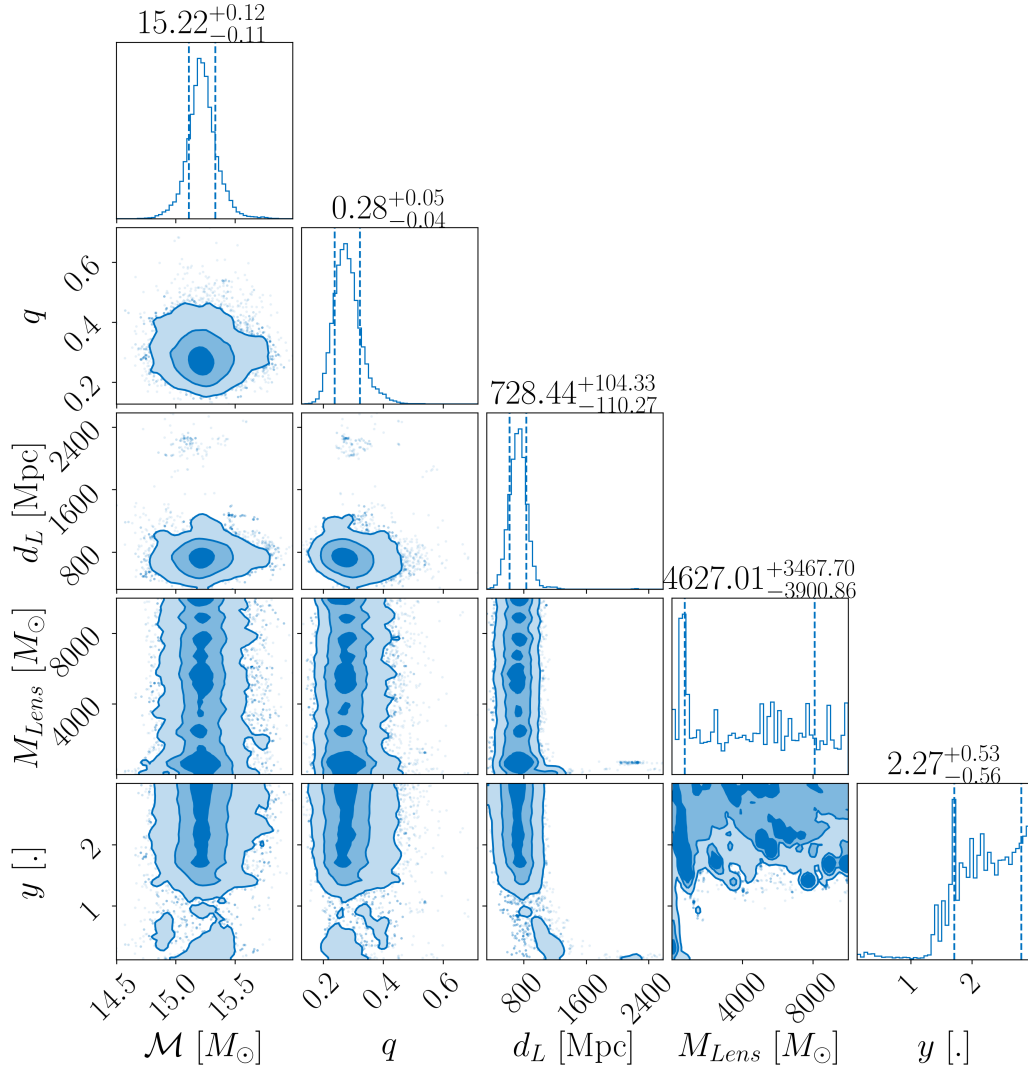


Figure 10.3.2: Posteriors for a subset of the parameters including the detector frame chirp mass and mass ratio, the luminosity distance, the lens frame lens mass, and the dimensionless displacement of the source from the optical axis (i.e. the source position). These posteriors were produced during the point mass microlensing analysis done for GW190412 using GRAVELAMPS. As can be seen, similarly to GW191103 (shown in Fig. 10.4.5), the lensing parameter posteriors are extremely broad and uninformative. This is consistent with the expectations for a non-microlensing event.

then passed on to GOLUM [198, 202] for filtering, and then to HANABI [235] for final analysis. The GW191103–GW191105 pair was identified as one of the most likely candidates by the PO analysis using the posteriors of some of the parameters obtained with the IMRPhenomXPHM waveform [295] released publicly on the Gravitational Wave Open Science Centre (GWOSC) [182, 349], whereas LENSID—which uses Q-transform images and BAYESTAR [319] skymaps—had not classified it as a

candidate. Appendix 10.A discusses possible reasons for the non-identification of the pair by the machine-learning based pipeline. For PO, the pair ended up having the highest overall statistic: $\log_{10} \mathcal{B}^{\text{overlap}} = 3.03$ and $\log_{10} \mathcal{R}^{\text{gal}} = 1.14$ for the SIS model giving a total of 4.17. Fig. 10.4.1 shows the posteriors of both events where one may see the varying degrees of overlap between the events.

Fig. 10.4.2 shows the candidate event pairs identified by PO analysis on the $\log \mathcal{B}^{\text{overlap}} - \log \mathcal{R}^{\text{gal}}$ plane considering both the SIS and SIE galaxy models. The choice of model affects only the $\log \mathcal{R}^{\text{gal}}$ value. The PO analysis is marginalised over phase and is, therefore, insensitive to the relative Morse phase ($\Delta\phi$) between the two events. As a result of this insensitivity, the SIE cases $\Delta\phi = 0$ and $\Delta\phi = \pi/2$ are considered separate models, hence we compute the \mathcal{R}^{gal} expected distributions for each case.

Posteriors of events detected by the LVK detectors can overlap by random coincidence meaning that unlensed pairs could also give high Bayes factors. For this reason, a background injection study with ~ 1000 unlensed events (the combinations of which yield about half a million pairs) is done to calculate the FAP [94] of the observed log Bayes factor for the candidate pair. The FAP per-pair (FAP_{pp}) for the candidate, hence the number of unlensed events with a Bayes factor higher than the one observed for the pair of interest, is found to be 1 in 10,000. Taking into consideration that a total of ~ 68 BBH events were detected in O3 the total FAP (given by $\text{FAP} = 1 - (1 - \text{FAP}_{\text{pp}})^{N_{\text{pairs}}}$) is found to be 0.3, i.e. a significance of slightly above 1σ . As seen in the figure the time delay for this event pair is more compatible with an SIE with $\Delta\phi = 0$ as compared to an SIE with $\Delta\phi = \pi/2$ and SIS. After this step, to extract more information about the event pair, it is passed to more extensive pipelines for further investigation.

Waveform Systematics Study on Posterior Overlap Candidate Identification

At the time of writing, no dedicated studies of waveform systematics have been conducted for gravitational lensing analyses. As an initial check, we report on a comparison of PO calculations on single-event PE performed with different waveforms. This is a practical first step as the single-event PE is significantly cheaper computationally than JPE, and detailed studies of waveform systematics on the latter are left for future work. The results presented here for the GW191103–GW191105 pair are an excerpt from a larger pioneer study on waveform model systematics in GW lensing that will be published separately [171].

The PO statistics reported in Abbott et al. [32] and used to initially qualify the

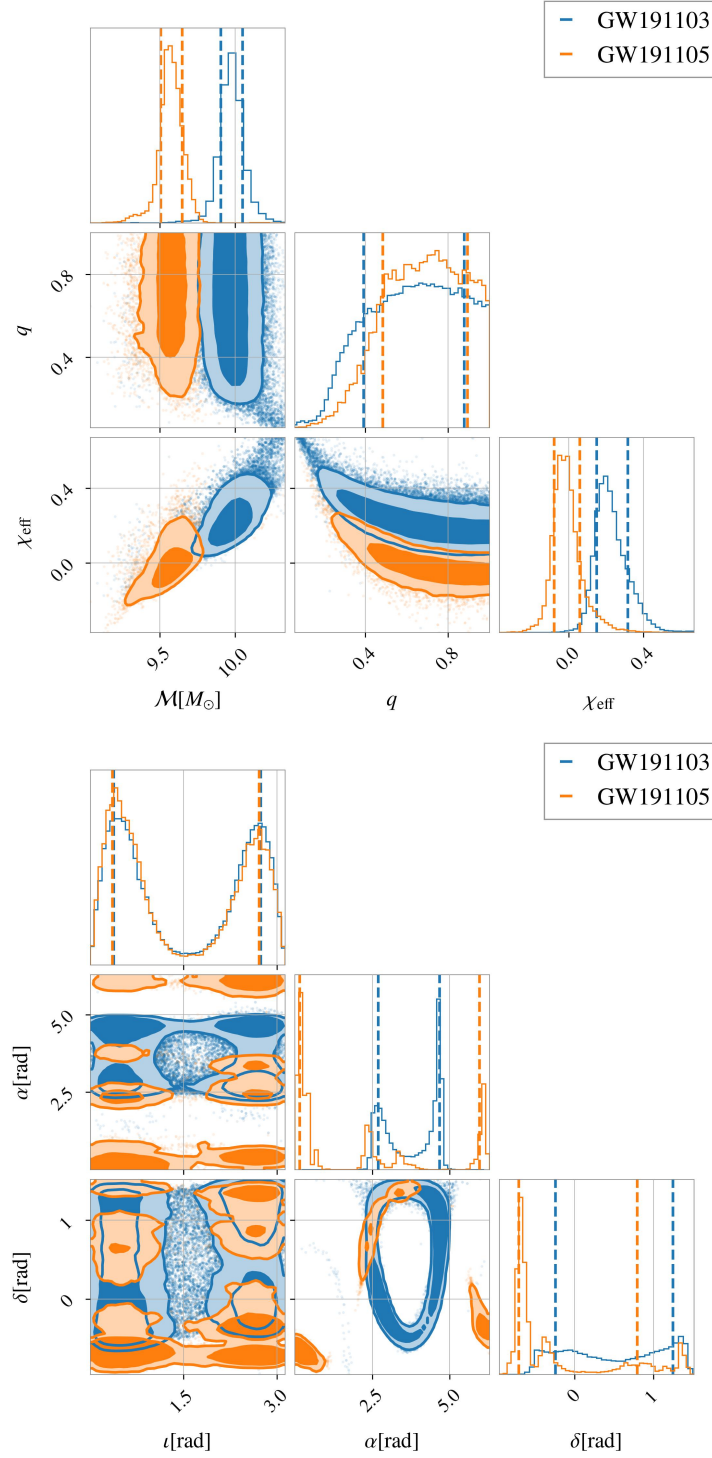


Figure 10.4.1: Posteriors for some of the parameters obtained using the IMRPhenomXPHM waveform for GW191103 (blue) and GW191105 (orange). The overlap in the extrinsic parameters (e.g. sky location) is larger than that for the intrinsic parameters (e.g. detector-frame chirp mass and spins).




Figure 10.4.2: The highest ranked candidate strong lensing pairs from the PO analysis considering all the event pairs found based on the O3 data (dots) [32] and the supra-sub pair analyzed in this work. The dashed lines correspond to the 1σ and 2σ confidence levels for the combined PO statistic ($\mathcal{B}^{\text{overlap}} \times \mathcal{R}^{\text{gal}}$) with different lensing models computed from the background simulations. We note that beside GW191103–GW191105, the pair analyzed in this work, GW190728–GW190930 is also close to 1σ for PO. However, the pair has been discarded in previous searches with a lower overall significance than GW191103–GW191105 [23]. Therefore, it is not considered for further analyzes in this work.

GW191103–GW191105 pair as sufficiently interesting for further follow-up were based on the IMRPhenomXPHM waveform [295]. Besides the posterior samples for this waveform, the data release [182] for GWTC-3 [30] contains samples obtained with the SEOBNRv4PHM waveform [285]. For this study, we performed additional PE runs on GW191103 and GW191105 for several other models, using PARAL-

LEL_BILBY [49, 326] with the DYNesty sampler [333], using settings and priors consistent with the GWTC-3 IMRPhenomXPHM runs [30].

The additional waveforms covered are three variants from the same family of frequency-domain phenomenological waveforms as IMRPhenomXPHM, as well as one time-domain phenomenological waveform:

- IMRPhenomXAS is an aligned-spin frequency-domain waveform for dominant-mode-only GW emission [294];
- IMRPhenomXHM is an aligned-spin frequency-domain waveform including HOMs [170];
- IMRPhenomXP is a frequency-domain waveform allowing for spin precession but for dominant-mode-only GW emission [295];
- IMRPhenomTPHM is a time-domain waveform allowing for spin precession and including HOMs [151].

The three reduced-physics IMRPhenomX waveforms allow us, in comparison with the most complete family member IMRPhenomXPHM, to check if neglecting any of these features has a significant impact on the PEs for each event, and hence on their overlap. In addition, the IMRPhenomTPHM waveform shares its time-domain nature with SEOBNRv4PHM but much of its modelling approach with IMRPhenomXPHM, making it an ideal tool to further cross-check for consistency between different modelling strategies.

We have followed the original KDE-based calculation from Haris et al. [188] to compute the PO statistic $\mathcal{B}^{\text{overlap}}$, with the modification of computing sky overlaps and intrinsic-parameter overlaps separately and then multiplying them, as done in Abbott et al. [32].

Table 10.4.1 lists the $\mathcal{B}^{\text{overlap}}$ resulting from comparing the posteriors from both events with each waveform. There are changes up to factors of ~ 6 in overlap statistics, with IMRPhenomXHM producing the highest $\mathcal{B}^{\text{overlap}}$ and SEOBNRv4PHM resulting in the lowest. We first notice that the aligned-spin waveforms produce the highest $\mathcal{B}^{\text{overlap}}$. These have fewer free parameters than the precessing models and hence also different prior volumes. By inspecting the posteriors we find that the aligned-spin waveforms prefer a_2 closely peaked to zero for both events which gives a high contribution to the overlap, while the precessing waveforms have broader

Table 10.4.1: PO statistic values for the GW191103–GW191105 pair using different waveform models in the single-event PE.

Waveform	$\log_{10}(\mathcal{B}^{\text{overlap}})$
IMRPhenomXAS	3.37
IMRPhenomXHM	3.48
IMRPhenomXP	3.08
IMRPhenomXPHM	3.03
IMRPhenomTPHM	2.70
SEOBNRv4PHM	2.65

distributions in this parameter, compensating with the additional freedom in the tilt angles. In addition, the two time-domain waveforms produce lower $\mathcal{B}^{\text{overlap}}$ than the frequency-domain waveforms. However, these changes are not significantly larger than expected from other sources such as prior choice and KDE implementation details, and all results are consistently in favour of the lensing hypothesis. Hence, we conclude that waveform choice does influence the PO method to some degree, but that for this specific event pair it is sufficiently robust under waveform choice, in the sense that all results agree qualitatively on identifying the pair as possibly lensed and interesting for follow-up.

One caveat on this type of study is that a full interpretation of $\mathcal{B}^{\text{overlap}}$ requires extensive simulation studies on both lensed and unlensed pairs, and the respective distributions could easily be different for different waveforms. However, in Abbott et al. [32], this factor was used purely as a ranking statistic. So, as long as the number does not drop strongly for any of the considered waveforms, we can conclude that the identification of the candidate pair is robust under waveform choice.

Compatibility with Lensing Models

Once an event has been identified as a potential candidate through the aforementioned PO or machine learning searches, it can be followed up by other pipelines. However, an additional check can be made by comparing the observed lensing parameters with the ones predicted by specific lensing models [188, 201, 235, 259, 364]. In this work, we focus on galaxy lenses. A comparison of the lensing parameters observed in the O3 search with the catalogue presented in More and More [259] and Wierda et al. [364] is given in Fig. 10.4.3. Most of the events are compatible with the values expected for unlensed events. Noticeably, two supra-

threshold event pairs have lensing parameters that are more consistent with a lensed hypothesis: GW191103–GW191105, which is the one the most on the left and therefore the most favoured, and GW190706–GW190719. One can also compute the \mathcal{M}^{gal} [259] statistics for the pairs. This number is the ratio of the probabilities for observing the lensing parameters under the lensed and unlensed hypotheses given by the lensing catalog though it does not account for the rate of lensing². For GW191103–GW191105, we find $\log_{10}(\mathcal{M}^{\text{gal}}) = 1.3$, while for GW190706–GW190719, $\log_{10}(\mathcal{M}^{\text{gal}}) = 0.8$. This shows that in the two cases, based on the catalogue, the observed lensing parameters agree more with the expected values for the lensed hypothesis than for the unlensed one.

Whilst such a comparison is valuable to gain information on the chances of lensing from specific models, it does not account for the compatibility of the binary parameters. Since the lensed hypothesis is favored over the unlensed one both based on PO and lensing statistics, we need to further ascertain the lensed nature by turning to JPE methods.

Joint Parameter Estimation Based Investigations

In the case of events being genuine lensed images, in addition to the lensing parameters being compatible with at least one lensing model, parameters whose estimation are unaffected by the lensing—e.g. sky location, component masses—should be the same between the events. For the GW190706–GW190719 pair, the GOLUM analysis gives $\log_{10}(C_U^L) = 0.6$, which is a value slightly favouring lensing, but still well within the values one can expect for unlensed events [201]. On the other hand, for GW191103–GW191105, we find $\log_{10}(C_U^L) = 2.6$. However, despite this higher value, this does not guarantee that the signals are not merely coincidentally similar [201]. Here, we follow up only on the GW191103–GW191105 pair, which has the highest coherence ratio observed in the O3 pairs [32], with the methodologies described in [201] and [235] to include information from a lens model into the analyses.

We use the method from Janquart et al. [201] to include a lens model in the detection statistic coming from the GOLUM pipeline [198, 202], modifying the coherence ratio so that it also accounts for the observed lensing parameters and a given lens model. This reduces the risks of false alarms in lensing searches and makes the detection of genuinely lensed pairs more robust. This also enables one to compare

²Such a catalogue is obtained through extensive lensed and unlensed populations simulations [188, 259, 364].

Table 10.4.2: Values of the detection statistic obtained using the GOLUM framework for the GW191103–GW191105 lensed candidate pair without lens model (C_U^L), and with an SIE lens, with ($C_{M_{\mu,t}}$) and without (C_{M_t}) relative magnification accounted for. The FAP_{PP} is decreased when using an SIE model.

Statistic	\log_{10} value	FAP_{PP}
C_U^L	2.5	2.0×10^{-3}
$C_{M_{\mu,t}}$	2.4	1.6×10^{-3}
C_{M_t}	2.9	9.8×10^{-4}

the compatibility of the observation with different lens models, constraining the nature of the lens.

To explore the event’s significance, we make an injection study by generating an unlensed background. We simulate 250 unlensed BBHs, sampling their masses from the `POWERLAW + PEAK` distributions; the spins and redshift are also sampled from the latest LVK observations, using the maximum likelihood parameters to generate the distributions [29]. The sky location is sampled from a uniform distribution over the sphere of the sky. The inclination is uniform in cosine, and the phase and polarizations are sampled uniformly on their domain. The merger times are set uniformly throughout the third observing run. For each set of parameters, we randomly associate a real event from the GWTC-3 catalogue and assume the same observation conditions (the same detectors are online, and the noise is generated from the event’s PSD). The 250 events selected are such that their network SNR is higher than 8³.

Based on this population, we can compute the FAP_{PP} for the coherence ratio for a given lens model. Table 10.4.2 summarizes the detection statistic and the FAP_{PP} for the analysis with and without model.

As a first step, we can verify the FAP_{PP} for the event pair when considering the coherence ratio. In this case, $\text{FAP}_{\text{PP}} = 2.007 \times 10^{-3}$, meaning that for our 250 unlensed BBHs, the chance is 1 in 500 that these events are not a genuine lensed pair. Thus, based only on the match between the parameters, the probability for these events to originate from two unrelated unlensed events is relatively high. Statistically, this means that the combination of only 33 randomly-selected unlensed

³ Whether an event can be considered to be of astrophysical origin is not dependent only on its SNR, and recent GWTC papers also put a threshold on the probability p_{astro} [30, 137]. Here, we consider the SNR threshold sufficient to assess detectability.

BBH mergers is capable to make at least one pair with at least the same coherence ratio as the GW191103–GW191105 pair.

Including a lens model helps study which astrophysical object could be at the origin of the lensed event. Here, we consider an SIE model [218, 259], which is a good model for a galaxy lens. We consider the case where we account for the time delay and the relative magnification in the model, and the case where we only consider the time delay. The combinations of the coherence ratio and the lensing statistics are written $C_{M_{\mu,t}}$ and C_{M_t} for the case with and without relative magnification respectively. The values for detection statistics and the FAP for the two cases are given in Table 10.4.2. The inclusion of the SIE model, with and without the relative magnification, decreases the FAP, meaning that the candidate pair becomes more significant. It is the case even if the SIE model with the time delay and the relative magnification slightly decreases the coherence ratio. The reduction in FAP is slightly larger when only the time delay is considered. This is because the observed relative magnification is slightly outside the highest-density regions for the two models under the lensed hypothesis. Moreover, the overlap between lensed and unlensed populations for this parameter is high, making it less helpful to discriminate between the two situations. The results have also been cross checked using an SIE-based catalog from Wierda et al. [364], giving the same conclusions.

We note here that whilst the FAP_{PP} seem relatively small for the SIE model, it is insufficient to claim the pair to be lensed. The smallest value found is 9.8×10^{-4} . Whilst this is an improvement over the original FAP_{PP} , it represents only an increase from the combinations of 33 randomly-selected unlensed BBH mergers to the combinations of 47 such mergers to, statistically, have a pair display a detection statistic higher or equal to the one reported for the observed pair. Consequently, whilst GW191103–GW191105 displays some interesting behaviours, these are insufficiently significant to claim a first strong lensing detection.

Additionally, we repeat the Bayes factor calculation comparing the probability ratio of the lensed versus the unlensed hypothesis as described in [32] using the more realistic lens population model described in [281] (see [236] and also Sec. 10.2) using HANABI [235]. We use the same set of source population models as in [32], e.g. the POWERLAW + PEAK model for the source masses from the GWTC-3 observations [29] and three models for the merger rate density: Madau-Dickinson [239], $\mathcal{R}_{\text{min}}(z)$, and $\mathcal{R}_{\text{max}}(z)$. Table 10.4.3 shows the log-10 Bayes factors computed using the three merger rate density models with the simple SIS lens model reported in [32] and the

Table 10.4.3: $\log_{10} \mathcal{B}_U^L$ for the GW191103–GW191105 pair from HANABI assuming three different merger rate density models and two different lens models. The values computed using the SIS model are reproduced from [32] for the sake of comparison. We see that the values with the SIE + external shear model are consistently higher than that with the SIS model, indicating a higher compatibility of the pair with a more realistic strong lensing model. However, since the values remain negative, the event is still most likely to be unlensed considering a more realistic lensing population with the most recent population models.

Lens model	Merger rate density	Madau-Dickinson	$\mathcal{R}_{\min}(z)$	$\mathcal{R}_{\max}(z)$
SIS [32]		−3.27	−3.21	−2.33
SIE + external shear [236]		−2.60	−2.46	−1.28

SIE + external shear model reported in [236]. We see that the values calculated using the SIE + external shear model are consistently higher than those using the SIS model, indicating that the pair is more consistent with a more realistic strong lensing model. Still, the $\log_{10} \mathcal{B}_U^L$ values are negative, and therefore the event pair is most likely unlensed.

Electromagnetic Follow-Up

In the case of a genuinely lensed GW signal, the light emitted by the host galaxy should also be lensed [187, 363]. As a result of this, in the case of a high-significance lensing candidate, one practical avenue would be to initiate an EM follow-up search. Identifying the host galaxy would be a way to verify the lensed nature of the signal independently.

The EM counterpart of a signal may be searched for in two ways: cross-matching of the joint GW sky localisation with EM catalogues containing known lens-source systems or a dedicated EM search on a per-event basis. Both of these make use of the improved GW sky localisation from the observation of multiple images which provide a virtually extended detector network [122, 198, 202, 233, 235]. A dedicated, per-event, EM search could be performed by looking for lenses within the sky localisation area and performing lens reconstruction to try to identify the specific lens at the origin of the observation [187, 363].

We cross-matched the GW191103–GW191105 pair with a few lens catalogues from optical surveys such as SuGOHI [97, 195, 196, 330–332, 368, 369], AGEL [345] and the Master Lens Database [MLD, 261]. Whilst no matches were found from the SuGOHI and AGEL catalogues, the grade A and grade B lenses ⁴ selected from

⁴ In this context, grade A lenses have a higher observational quality and accuracy than grade B

MLD at galaxy scales showed 5 matches—4 doubly lensed systems and 1 quadruply lensed system (see Fig. 10.4.4). Two of the doubles are predicted to have time delays > 50 days based on the best-fit lens mass models [204, 290]. For the remaining double, we infer a time delay of ~ 20 days given the redshifts of the lens galaxy and the source as well as the velocity dispersion [73] under the assumption of an SIS lens mass distribution. All of these time delays are too long to be consistent with the 2-day time gap of the GW191103–GW191105 pair.

Lastly, we determine the time delays expected for the quadruplet SDSS J0918+5104 using the best-fit mass model and results from Ritondale et al. [301]. The expected time delay for the closest pair in SDSS J0918+5104 is around $\sim 0.5 - 1$ day. Given the uncertainties in the lens model, this ballpark estimate of time delay could be consistent with that of the GW191103–GW191105 pair. However, a more detailed mass modelling analysis and exploring different physical scenarios for the offset between the host galaxy and the GW source can lead to slightly different degrees of compatibility. Still, the predicted relative magnification is both qualitatively and quantitatively consistent since the latter GW event is found to be weaker than the former GW event according to the HANABI analysis. Furthermore, the closest pair of SDSS J0918+5104 images corresponds to a minimum (Type I) and a saddle point (Type II) suggesting a phase shift of $\Delta\phi = \pi/2$. This is somewhat less favored but yet plausible for the GW191103–GW191105 pair based on the PO analysis (see Fig. 10.4.2).

We note that the cross-matching analysis is limited by both the incompleteness of the databases of known EM lenses and the algorithms used to find matching lenses. A more detailed investigation (see Wempe et al. [363] for an example of how to investigate the link between EM and GW lensed systems) is warranted to assess the probability of a lens like SDSS J0918+5104 to be a genuine EM counterpart of the GW191103–GW191105 pair, if the latter is considered to be lensed. In the future, having dedicated EM follow-up of the lensed event candidates using optical telescopes could also help to gather more information about the lens and localise the source to the host galaxy.

Microensing Analysis

Whilst interest in the GW191103–GW191105 pair was triggered from the strong lensing perspective, it is worth paying additional attention to whether either of the

ones.

single events in the pair displays any signs of frequency-dependent effects associated with microlensing. As has been noted above, the most likely microlensing scenario is a microlens embedded within a macrolens. In such a scenario, one or both of the individual signals could display the signatures of the microlens [257, 316]. To determine if that scenario is plausible both GW191103 and GW191105 have been analysed using the GRAVELAMPS pipeline to determine the evidence for an isolated point mass or SIS microlens.

The main result of this analysis is that neither event shows particular favouring for either the point mass or the SIS microlensing models over the unlensed model. The preferred case is that of an SIS microlens in GW191103 which produces a $\log_{10}(\mathcal{B}_U^L) = 0.38$. In the point lens case, however, support drops to a $\log_{10}(\mathcal{B}_U^L)$ of 0.09 meaning that neither case posits strong evidence. To further compound this, the posteriors for the SIS case are given in Fig. 10.4.5. The source parameters' posteriors are well constrained, but those of the lensing parameters are extremely broad and uninformative. This is consistent with the expectation of an unlensed event and, combined with the marginal Bayes factor, leads to the conclusion that there are no observable microlensing signatures within this event.

The case of GW191105 is similar, albeit with even less favouring for the lensing models. Here, $\log_{10}(\mathcal{B}_U^L) = 0.21$ for the SIS case, and it is -0.35 for the point mass lens model. With even the more optimistic of the two models having a lower favouring, as well as a repetition of the posterior behaviour for the lensing parameters, one again concludes that there are no observable microlensing signatures within this event either.

Targeted Sub-threshold Search

Whilst lensing may produce multiple images, it is not guaranteed that all of the images will be detected. However, if it can be ascertained that a detected signal (or signal pair) is lensed, this allows deeper investigation for events below the detection threshold used for standard searches. Reciprocally, finding a sub-threshold counterpart to images with a low probability of being lensed could increase the support for the lensing hypothesis. As such, we conducted searches for sub-threshold lensed counterparts to GW191105 and GW191103 individually in the O3 data with the event-specific template banks constructed out of the intrinsic parameter posterior samples [223]. These searches yielded 7 triggers for GW191105 and 15 triggers for GW191103 above the false-alarm rate (FAR) threshold of 1 in 69 years as defined

in Abbott et al. [32]. None of these were reported as a potential lensed counterpart to any of the GW191103 and GW191105 events. One of the interesting triggers found was LGW191106_200820 which arrived just about a day after GW191105, agreeing with galactic scale lens models. However, this trigger was ruled out as a lensed counterpart to a GW191103–GW191105 pair as the overlap in the sky location is poor and the evidence for the event being a real event is very small. It was thus concluded that no promising candidates for an additional sub-threshold counterpart image for GW191103–GW191105 was found within the O3 data.

10.5 GW191230_180458–LGW200104_180425

During the O3 sub-threshold lensing counterpart search, the TESLA pipeline [223] based on the GstLAL software [86, 250] found roughly 470 triggers which could be potential strong lensing counterparts to the superthreshold events. Of these, two had a FAR lower than 1 in 69 years [32] though none were found to have support for the lensing hypothesis and all were ultimately discarded. An alternative method for identifying the sub-threshold triggers as possible lensed counterparts to superthreshold events, developed in Goyal et al. [180], uses the BAYESTAR localisation skymaps, matched-filter chirp mass estimates and the time delay priors to rank all the super-sub pairs. It identifies the sub-threshold event termed LGW200104_180425⁵ as a possible lensed counterpart to the supra-threshold GW191230_180458 event. It is the most promising super-sub pair according to this method as it has significant sky and mass overlap, coupled with apparent lensing parameters matching the expected values for a galaxy lensed models (see Goyal et al. [180] for more detail). In the rest of this section, we denote the super- and the sub-threshold events GW191230 and LGW200104, respectively.

LGW200104 was detected with both LIGO detectors with an SNR of 6.31 in Hanford and 4.94 in Livingston. The GstLAL matched-filter estimates on its chirp mass place it at $67.39M_{\odot}$ with the individual component masses being $82.48M_{\odot}$ and $72.71M_{\odot}$. These high component masses combined with the faintness of the signal contribute to a very low p_{astro} of 0.01 during usual unlensed super-threshold searches, where the event was found with the SPIIR [106, 237] and cWB [217] pipelines, signifying a significant lack of probability of the event being a genuine detection. Likewise, the FAR found for this event during the super-threshold searches is 4824/yr, also

⁵ Here, we follow the usual naming convention, adding an L at the start of the event name to specify it is a sub-threshold candidate. Therefore the name of the sub-threshold trigger is LGWYYMMDD_hhmmss, where YY is the year, MM the month, DD the day, hh the hour, mm the minutes and ss the second in UTC time.

favouring a terrestrial origin for the signal [206]. Since the sub-threshold searches have a more focused template bank, they also reduce the FAR for the events when they are in the correct region of the parameter space [223, 247]. Therefore, the FAR for the event decreases to 6.59/yr when it is found with the TESLA pipeline [223], still higher than the threshold used for following-up on sub-threshold events in O3 [32]. In keeping with the analyses done within this work, whilst we do not claim that the event is both genuine and genuinely lensed, we treat it as though it were. Consequently, we investigate the pair using the lensing identification tools used for super-threshold pairs.

PyCBC Sub-Threshold Search

To further verify this candidate and look for sub-threshold counterparts, an independent search pipeline, based on PyCBC [127, 351], has been applied [247]. In contrast to the GstLAL-based TESLA pipeline [223], the PyCBC-based approach uses a single template based on the posterior distribution of each target event. This search method was previously applied to O1–O2 data [247] and O3a data [23]. Whilst this search was not deployed across the totality of the O3b data, we have applied it as a cross-check on the chunk of data containing LGW200104 starting on 2019/12/03 at 15:47:10, and ending on 2020/01/13 at 10:28:01, looking for counterparts to GW191230.

For the template, we selected the maximum-posterior point from the IMRPhenomXPHM samples for GW191230 released with GWTC-3 [182], from a KDE after removing transverse spins, as to obtain the parameters for a single IMRPhenomXAS template with the following values: $m_1 = 82.48M_\odot$, $m_2 = 72.71M_\odot$, $a_1 = -0.0037$, and $a_2 = 0.026$. After running PyCBC over the chunk using the same clustering steps as in McIsaac et al. [247], the results are collected and events are ranked by the inverse of their false alarm rate (IFAR).

For the examined chunk, we found 5 candidates above an IFAR threshold of 1 month, with 2 previously known GW events topping the list, one being GW191230 itself. To check the correlation of the remaining 3 events with the target supra-threshold event, we performed a sky overlap estimation of each pair, following the idea described in [366]. The results are shown in Table 10.5.1. The sky overlap is computed as the fractional overlap between the sky map obtained using parameter estimation for GW191230 and the sky map produced using BAYESTAR [319] for each sub-threshold candidate. Since these two methods do not match exactly, it leads to a 75% overlap

Table 10.5.1: PyCBC targeted sub-threshold results for counterpart candidates to GW191230 ranked by IFAR. From left to right, the columns represent the event, the time delay compared to the supra-threshold event used to make the template, the inverse false-alarm rate (IFAR), the signal-to-noise ratio (SNR), and the 90% confidence region (CR) overlap for the sky posteriors.

Rank	Name	Event	ΔT [days]	IFAR [yr]	SNR	90% CR Overlap
1	LGW191222_033537	GW191230	8.60	125822.11	10.99	0.00
2	LGW191230	GW191230	0.00	312.15	10.11	0.75
3	LGW191212_220841	GW191230	17.83	0.57	16.38	0.00
4	LGW191214_055524	GW191230	16.51	0.10	7.16	0.02
5	LGW200104	GW191230	5.02	0.09	8.02	0.62

for the supra-threshold event with itself.

In interpreting sub-threshold search results, one has to take into account that there is a good chance that, in addition to the potential counterpart images, there will be candidates originating from instrumental glitches or also from different, weaker, GW events that were not identified in previous searches. Here, the candidate corresponding to LGW200104 is ranked fifth (including GWTC events) with an IFAR of 0.09 years. Its network SNR is recovered as 8.02 (with an SNR of 6.31 and 4.94 in H1 and L1, respectively) and its sky localization overlap with GW191230 is 62%. The sky overlap map is given in Fig. 10.5.2.

The third-ranked event has a higher SNR, but no sky overlap with GW191230 and can be clearly identified as a glitch since there is simply an excess in power for all frequencies at a given time and no time-frequency evolution similar to that expected for a CBC signal. The fourth-ranked event is clearly a case of a scattered-light glitch [327, 328, 343]. Appendix 10.B further details these two events. In the end, the PyCBC sub-threshold search also finds LGW200104 as the most plausible lensed GW sub-threshold counterpart to GW191230 consistent with the GstLAL-based TESLA pipeline [223] and the ranking method proposed in Goyal et al. [180].

Posterior Overlap Analyses

From the PO analysis this pair has $\log_{10} \mathcal{B}^{\text{overlap}} = 2.45$. Since the combined SNR of the sub-threshold trigger is close to 8, it is reasonable to treat the event pair the same way we did for other candidates. Using the same time delay priors as for the supra-threshold events we find $\log_{10} \mathcal{R}^{\text{gal}} = 0.97$ which makes the log of the overall

Table 10.5.2: Posterior-overlap factors for the GW191230–LGW200104 pair using different waveform models in the single-event PE.

Waveform	$\log_{10}(\mathcal{B}^{\text{overlap}})$
IMRPhenomXAS	3.20
IMRPhenomXHM	3.13
IMRPhenomXP	2.52
IMRPhenomXPHM	2.45
IMRPhenomTPHM	2.55

PO statistic 3.43. Fig. 10.5.3 shows the posteriors for LGW200104 and GW191230. Visually, the degree of overlap in both extrinsic and intrinsic parameters is high. However the intrinsic parameters posteriors are broader as compared to GW191103–GW191105. For events having high masses in the detector frame, such as these, the number of cycles in the waveform within the LIGO–Virgo frequency band is smaller. This leads to broader posteriors which in turn reduce the overlap statistics, while increasing the rate of coincidental overlaps [118, 162]. In addition, lensed events are more likely to have higher detector frame masses than unlensed events due to the their magnification. Hence, it is a challenge to identify high-mass lensed candidates. Including the population priors and selection effects might help [188, 201, 235].

We also compute the significance of the pair using the supra-threshold background introduced in Sec. 10.4 and find it to be $\lesssim 1\sigma$, as shown in Fig. 10.4.2. This implies that this pair, though not conclusively lensed, is one of the most significant candidates amongst all the O3 event pairs.

To look for potential waveform systematics, we perform the same analysis as in Sec. 10.4 using results from `PARALLEL_BILBY` runs with the same five waveform models in the posterior overlap calculation. The results are shown in Table 10.5.2 and we find relatively consistent results. However, we again notice that the aligned-spin waveforms produce higher $\mathcal{B}^{\text{overlap}}$, by a factor of ~ 6 . In this case, both a_1 and a_2 peak towards zero for the aligned-spin models for the two events, leading to a better overlap. However, all waveforms agree on identifying this pair as consistent with lensing.

The PO analysis can quickly identify the lensed candidates but it does not take into account the full correlation between the data streams, the selection effects, and the lensing parameters. Hence, the candidates are passed on to JPE pipelines for further

investigations.

Joint Parameter Estimation Based Investigation

After discovering the candidate with the sub-threshold searches and confirming interest with PO, it was analysed in more detail using GOLUM [198, 202] with the samples of the supra-threshold event as the prior for the sub-threshold one. The evidence of this run can be compared to the results of a standard unlensed BILBY [49] investigation to yield the coherence ratio. In this case, the run yielded $\log_{10}(C_U^L) = 1.1$. This is lower than that calculated for GW191103–GW191105. However, in this case, one of the two images is very close to the limit of a detectable event and this may impact the coherence ratio. By itself, the coherence ratio also is still high enough to favour the lensing hypothesis. To initially investigate the pair’s significance, it was compared with the same background as outlined in Sec. 10.4. This results in a FAP_{pp} of 1.4% and thus a FAP of 0.6 which indicates the event is consistent with a coincidental unlensed background event. However, the background resulting in this FAP consisted entirely of supra-threshold events and the exact effects of sub-threshold events in such studies have not been deeply explored.

The GOLUM analysis also offers the possibility to gain insight into the lensing parameters. In particular, it gives information about the difference in Morse factor and relative magnification ⁶. Their posterior distributions are given in Fig. 10.5.4 and 10.5.5, respectively, in which it can be seen that the relative magnification peak is at ~ 1.5 , meaning that its value is close to the highest-probability region expected for an SIE lens model (see for example Fig. 10.4.3). On the other hand, the difference in Morse factor is less well constrained, with the main support being for $\Delta n = 0.5$, but also some support for the $\Delta n = 0$ case. We note that, generally, for well-detectable lensing events, the difference in Morse factor is well recovered [198, 202]. This observation may indicate that the event is unlensed but also simply that the lower SNR of the signal makes the identification harder. These lensing parameters and the time delay, however, are consistent with expected values for a galaxy-scale lens.

Based on the GOLUM results, we may also investigate how the coherence ratio and the FAP evolve with the inclusion of expected parameter values from a lens model, as was done in Sec. 10.4. Using the same background, and the same models as within that section, we compute the population-reweighted coherence ratios. These values are reported in Table 10.5.3. Notably, the coherence ratio

⁶ It also gives the possibility to constrain the time delay, but since the arrival times are very well measured already in GW data analysis, this does not provide much additional information.

Table 10.5.3: Values of the detection statistic for the GW191230–LGW200104 lensed candidate pair without lens model (C_U^L), and with an SIE lens, with ($C_{\mathcal{M}_{\mu,t}}$) and without ($C_{\mathcal{M}_t}$) relative magnification accounted for. The FAP_{PP} is decreased when using an SIE model, showing that the observed characteristics are in line with the expected behaviour for the given model and population.

Statistic	\log_{10} value	FAP_{PP}
C_U^L	1.105	1.401×10^{-2}
$C_{\mathcal{M}_{\mu,t}}$	3.427	1.167×10^{-3}
$C_{\mathcal{M}_t}$	1.915	2.017×10^{-3}

found for the SIE model including both the relative magnification and the time delay is now higher than that for the GW191103–GW191105 event pair, meaning that the observed characteristics are more in agreement with the expected value for a galaxy-lens model and the currently observed population than for that pair. This is a demonstration of the fact that the candidate pair—even though the sub-threshold event is not confirmed to be of astrophysical origin—is interesting.

The GW191230–LGW200104 pair was also analyzed by the full JPE package HANABI [235] where the joint parameter space of the two events was simultaneously explored by the stochastic sampler DYNESTY [333] with settings identical to those used in [32]. The parameters found are consistent with the ones found using the GOLUM framework—see Fig. 10.5.5 for a comparison of the relative magnifications. In particular Fig. 10.5.6 shows the posterior probability mass function for the possible image types of the GW191230–LGW200104 pair. We see that the image type configurations for the two events that have non-zero support have the difference in the Morse phase factor Δn either 0 (i.e. the I-I, II-II and III-III configuration) or 0.5 (i.e. the II-I and III-II configuration). Again, this is consistent with the values shown in Fig. 10.5.4 obtained using GOLUM.

We also performed the Bayes factor calculation comparing the probability ratio of the lensed versus the unlensed hypothesis for this pair in the same fashion that we did for the GW191103–GW191105 pair as in Sec. 10.4. Again, we use the same set of source population models as in [32], e.g. the POWERLAW + PEAK model for the source masses from the GWTC-3 observations [29] and three models for the merger rate density: Madau-Dickinson [239], $\mathcal{R}_{min}(z)$, and $\mathcal{R}_{max}(z)$. Table 10.5.4 shows the \log_{10} Bayes factors computed using the three merger rate density models with the simple SIS lens model [32] and the SIE + external shear model [236]. We see

Table 10.5.4: $\log_{10} \mathcal{B}_U^L$ for the GW191230 and LGW200104 pair from HANABI assuming three different (source) merger rate density models and two different lens models. We see that the values with the SIE + external shear model are all positive (but only mildly) and consistently higher than that with the SIS model which are all negative, indicating a higher compatibility of the pair with a more realistic strong lensing model. Note that the calculations assumed that both GW events are astrophysical of origin. These values are not sufficient to claim the event pair to be lensed as we would require a positive \log_{10} posterior odds, and the observed Bayes factors are not high enough to balance the low prior odds for strong lensing.

Lens model	Merger rate density	Madau-Dickinson	$\mathcal{R}_{min}(z)$	$\mathcal{R}_{max}(z)$
SIS		−0.76	−0.35	−0.57
SIE + external shear		0.14	0.57	0.30

that the values calculated using the SIE + external shear model are positive but only mildly (< 1), and they are also consistently higher than the values computed using the SIS model (which are all negative), indicating that the pair is more consistent with a more realistic strong lensing model. It should be noted that the calculations assumed that both GW events are astrophysical of origin and the second is treated as a super-threshold event.

Despite some of the evidence for this event aligning relatively well with the expectations for a lensed event, there remain several key arguments against a claim of lensing for this pair. The first is that whilst it is the case that the event has the highest currently observed Bayes factor, it is insufficient to yield a positive log posterior odds considering that the \log_{10} prior odds is between -2 and -4 [81, 224, 266, 274, 281, 364]. The second argument, is the nature of the trigger itself. The sub-threshold event is not convincing—consider for instance the extremely low p_{astro} and FAR—and there is no clear evidence to claim that the event is a genuine GW detection.

In the end, although the event pair is unlikely to be lensed, the analyses performed on this event pair serve as a powerful demonstration of the necessity for searching for such sub-threshold counterparts and the kinds of information that they may yield.

Electromagnetic Follow-Up

Even though the lensing hypothesis is disfavored, we investigate if there are any EM lens systems with consistent lensing properties from the literature for this event pair. As in Sec. 10.4, we cross-matched with the MLD. The grade A and grade B lenses

selected from the catalog at galaxy scales showed 21 matches (see Fig. 10.5.7). There are two major lens samples that fall within the 90% confidence region of the sky localisation in addition to a handful of systems from heterogeneous studies [e.g., 318, 346]. The Strong Lensing Legacy Survey (SL2S) lens systems are those with $RA < 40 \text{ deg}^2$ [172, 260, 329] and South Pole Telescope (SPT)–ALMA lens systems are those with $RA > 50 \text{ deg}^2$ [e.g., 334, 362]. Some of these lens systems do not have sufficient information (for instance, are lacking source redshift) and the others do not have best-fit mass model parameters or sufficiently high-resolution imaging to identify the multiple lensed image positions. In the absence of these constraints, the time delay of a few days (i.e. ~ 5 days for the GW191230–LGW200104 pair) can easily be consistent with many of the lens systems. As a result, a rudimentary analysis to determine the likelihood of these lens systems and to select the most likely EM counterparts is not possible, and a more detailed mass modelling for all of the 21 systems would be necessary to find the most promising EM counterpart. Whilst the observation of a third or a fourth GW image would also help further in constraining the characteristics, the lack of high resolution imaging to clearly and accurately measure the image positions are still anticipated to be the limiting factor. Lastly, we again caution the reader that any incompleteness of the survey (both telescope imaging and subsequent lens searches) may mean that additional potential EM counterparts may have been missed from our initial list of 21 candidates. We found more candidates as compared to the EM follow up of GW191103–GW191105 (Sec. 10.4) merely because there are more optical surveys that have looked towards the sky region of interest here with respect to the sky overlap region of GW191103–GW191105 which is nearer to the poles. Hence in order to have a robust association one needs to incorporate the selection effects for both the EM and GW observations (see [363] for possible avenues).

10.6 GW200208_130117

GW200208_130117, denoted GW200208 from here on, was selected for follow-up in this paper for two reasons. The first was because it was the event with the highest $\log_{10} \mathcal{B}_U^L$ in the O3 microlensing analysis [32], with a value of 0.8 which, whilst positive, is still inconclusive. The secondary reason was that the event had a relatively narrow posterior on the redshifted lens mass which is atypical of unlensed events. In the O3 lensing paper it was considered that the cause of the apparent favouring of microlensing for the event could be due to short duration noise fluctuations causing an apparent dip in the signal, mimicking the beating pattern of

microlensing [32].

Microlensing Model Investigation

As has been done with a selection of the other events within this paper, GW200208 was re-examined using the GRAVELAMPS pipeline [370] to investigate the potentiality of model selection in the case of a microlensing candidate. Whilst the isolated point mass model used by GRAVELAMPS is similar to that used by the O3 microlensing search pipeline, there are sufficient implementation differences to warrant re-examination with GRAVELAMPS for this model rather than simply comparing the results of the SIS investigation with those of the O3 microlensing analysis pipeline.

For all of the models examined, GRAVELAMPS had increased favouring for the microlensing hypothesis with this event as compared to the O3 microlensing analysis pipeline. In the point lens case, the $\log_{10} \mathcal{B}_U^L$ increases to 1.20. This confirms the result from the O3 analysis and shows the event warrants additional investigations. In the SIS case, the preference increased further with the $\log_{10} \mathcal{B}_U^L = 1.77$. This again is sufficiently high as to warrant additional scrutiny, but not high enough to make a statement by itself.

One stage of preparatory work that would shed additional light on the potential significance of the $\log_{10} \mathcal{B}_U^L$ figures would be a detailed background study to determine the range over which unlensed events may appear as microlensing candidates. Such a study had been done for the microlensing search in [23, 32] and allowed the statement above that for the case of that pipeline, the \mathcal{B}_U^L recorded for GW200208 was within the expected range for an unlensed event. Due to computational constraints, this has not yet been performed for the GRAVELAMPS models, although it is one of the steps that should be taken during O4 so that the significance of candidates may be evaluated quickly. What we note is that the increase between the two pipelines would not necessarily have rendered GW200208 outside of expectations for the O3 microlensing search pipeline, and the general trend of events analysed would appear to indicate that in general SIS is preferred over the point mass model. This is likely due to a lens with similar parameters producing lower peak amplifications in the SIS model as compared with the point mass model which would yield smaller deformations from the unlensed template.

For the other events that have been examined, the posteriors for lensing parameters have been a factor in determining that the microlensing hypothesis is unlikely. However, in the case of GW200208—the posteriors of which in the SIS case are

Table 10.6.1: This table presents the results of a Bayesian model comparison study between the unlensed and the microlensed hypotheses for GW200208, with microlensing model corresponding to an isolated point-lens mass. The study was conducted for different configurations and sampler settings, as indicated by the different columns, to verify for possible noise artifacts and check the influence of the sampler settings on the results. The table includes the waveform approximant used, the lower and higher frequency cutoffs used for likelihood evaluation (f_{low} , f_{high}), duration of the data segment used, and the priors on the redshifted microlens mass (M_{Lz}^z) and the impact parameter (y), represented by $p(M_{Lz}^z)$ and $p(y)$ respectively. The Bayes factor for the support of microlensing over the unlensed waveform model is given by $\log_{10} \mathcal{B}_U^L$. The range of the priors is also indicated. The terms ‘L.U’ and ‘L.L.U’ refer to log-normal and log-log-normal distributions respectively, while ‘P.L’ refers to a power law profile with the index given by α .

Waveform	f_{low}	f_{high}	duration	$p(M_{Lz}^z)$ [min,max]	$p(y)$ [(α), min, max]	$\log_{10} \mathcal{B}_U^L$
IMRPhenomXPHM	20	448	4	L.U [1, 10^5]	P.L [1, 0.1, 3.0]	0.89
IMRPhenomXPHM	20	1024	4	L.U [10, 10^5]	P.L [1, 0.01, 5.00]	0.63
IMRPhenomXPHM	20	896	8	L.U [10, 10^5]	P.L [1, 0.01, 5.00]	0.46
IMRPhenomXPHM	15	448	4	L.U [10, 10^5]	P.L [1, 0.1, 3.0]	1.02
IMRPhenomXPHM	15	448	4	L.U [10, 10^5]	P.L [1, 0.01, 5.00]	0.53
IMRPhenomXPHM	15	448	4	L.U [10, 10^5]	Uniform [0.1, 3.0]	1.04
IMRPhenomXPHM	15	448	4	L.L.U [10, 10^5]	P.L [1, 0.1, 3.0]	0.70
IMRPhenomXPHM	15	448	4	L.L.U [10, 10^5]	Uniform [0.1, 3.0]	0.95
IMRPhenomXPHM	15	448	4	Uniform [10, 10^5]	Uniform [0.1, 3.0]	0.50
NRSur7dq4	20	448	4	L.U [1, 10^5]	P.L [1, 0.1, 3.0]	0.96
NRSur7dq4	18	448	4	L.U [1, 10^5]	P.L [1, 0.1, 3.0]	0.90

shown in Fig. 10.6.1—this same test yields results more consistent with the microlensing hypothesis. As can be seen in the figure, the lensing parameter posteriors are relatively narrowly constrained around a $2000M_\odot$ lensing object with a source position value of 0.60. Fig. 10.6.2 shows that in the more pessimistic point lens case, the lensing parameters are constrained to similar values which further cements the need for additional scrutiny of this event. For the two models, we see that the $3\text{-}\sigma$ confidence intervals for the lensing parameters are a bit noisy. However, the peaks in the density distributions remain clearly visible.

We further investigate GW200208 using various sampler and prior settings, as well as testing different waveform models, as listed in Table 10.6.1. These tests are designed to verify whether noise artifacts could be at the root of the observed support

for microlensing, and whether the results are robust for different sampler settings. We assume the microlensing model of an isolated point-lens and do parameter estimation using the GWMAT framework [255].

This pipeline utilizes a C_{YTHON} implementation of the amplification factor calculation for the isolated point mass model serves as an independent cross-check for the study. Additionally, the pipeline incorporates a dynamic cutoff based on the source position y to transition to a geometric optics approximation. The resulting probability density functions representing the recovered microlensing parameters are illustrated in Fig. 10.6.3.

Firstly, we observe that the posteriors for both parameters, y and $\log_{10} M_{\text{Lens}}^z$, are consistent across different runs, with median values and 1-sigma errors of $1.07^{+0.61}_{-0.32}$ and $3.15^{+0.18}_{-0.21}$, respectively. However, the posteriors for y show signs of railings at the upper end, particularly in runs with an upper limit of 5 in the prior. Since the SNR is around 11, we limit the prior on y to a maximum of 5, which would already require an $\text{SNR} \gtrsim 40$ to make the microlensing signatures detectable ⁷. It is worth noting that the posteriors for $\log_{10} M_{\text{Lz}}$ are relatively well-converged, with a sharp peak, except in the case where we used a uniform prior in M_{Lens}^z , which shows a tendency towards bimodality with another peak at $\log_{10} M_{\text{Lens}}^z \gtrsim 5$.

As shown in Table 10.6.1, we primarily use $f_{\text{high}} = 448 \text{ Hz}$ ⁸ and a duration of 4 seconds due to the heavy mass nature of the GW200208 event, which has a total binary mass of approximately $90 M_{\odot}$ and negligible spin content. Comparing the first entry in the table with the second-to-last entry, we find that NRSur7dq4 [354] yields a slightly higher Bayes factor value than IMRPhenomXPHM for similar settings, except that the NRSur7dq4 case imposes a total mass constraint of greater than $66 M_{\odot}$, considering its region of validity for $f_{\text{low}} = 20 \text{ Hz}$. However, since the event has a total mass $\geq 74 M_{\odot}$ with 3σ certainty, we also analyzed the event with $f_{\text{low}} = 18 \text{ Hz}$ (last row), resulting in a slight decrease in $\log_{10} \mathcal{B}_{\text{U}}^{\text{L}}$. On the other hand, when we lowered the value of f_{low} to 15 Hz for the IMRPhenomXPHM case, the $\log_{10} \mathcal{B}_{\text{U}}^{\text{L}}$ increased compared to a similar run with $f_{\text{low}} = 20 \text{ Hz}$ (see 1st and 4th entry).

When we choose $f_{\text{high}} = 896 \text{ Hz}$ and a duration of 8 s (3rd row), both $\mathcal{B}_{\text{noise}}^{\text{U}}$ and $\mathcal{B}_{\text{noise}}^{\text{L}}$ decrease, as does $\log_{10} \mathcal{B}_{\text{U}}^{\text{L}}$, resulting in the lowest value among all the different

⁷ The minimum SNR required to distinguish two waveforms with a match value of m is roughly $\sqrt{2/(1-m^2)}$, where we used an estimate of the Bayes factor and set a threshold of unity for distinguishability [113, 135, 353]. For $M_{\text{Lens}}^z = \mathcal{O}(10^3 M_{\odot})$, the match value between the unlensed and lensed waveform with $y = 5$ comes out to be ~ 0.9993 , which implies a minimum SNR of ~ 40 .

⁸ We choose $f_{\text{high}} \leq 0.875 * (f_s/2)$, where f_s is the sampling frequency (see Appendix E of [30]).

settings used in the table. Additionally, $\log_{10} \mathcal{B}_U^L$ also decreases when we broaden the prior on y or $\log_{10} M_{Lz}$ (compare, for example, the 4th and 5th row), which could be additionally lowered due to railing and bimodalities at the higher values of y and $\log_{10} M_{\text{Lens}}^z$, as shown in Fig. 10.6.3.

The apparent railing and the bimodality can be attributed to the fact that if the likelihood fails to exhibit strong unimodality, the posterior densities may vary depending upon the prior beliefs. A higher upper limit in the prior of y with a power law profile $p(y) \propto y$ will assign more weight to higher values of y . Similarly a uniform prior in M_{Lens}^z places a higher weight on heavier microlenses than a log-uniform or log-log-uniform prior, thereby increasing the posterior density in that region. However, if the SNR is high, or if the event is truly microlensed, the likelihood values are better constrained and the posterior densities would not be expected to change much with the priors.

We also note that the Bayes factors presented in Table 10.6.1 show more variability. These results indicate that we cannot make a firm conclusion on whether the event is microlensed or not based on the Bayes factor, and the event can only be deemed interesting probabilistically depending on the prior beliefs we choose.

Maximum Likelihood Injection

One avenue of investigation to determine whether an event with the parameters that are suggested by the lensing models within GRAVELAMPS would be detected, and if it was detected, how significant a detection would we expect is to examine a simulated waveform with the maximum likelihood parameters injected into simulated detector noise.

Whilst as stated above, a full-scale injection campaign was not undertaken for the GRAVELAMPS analysis due to temporal and computational constraints, we can investigate if the \mathcal{B}_U^L figures would be plausible for a genuine microlensing event of the suggested parameters by injecting a signal with the parameters of the maximum likelihood sample of the GRAVELAMPS analysis into a realisation of Gaussian noise assuming a representative PSD for the noise around the time of detection and analysing this injection with the GRAVELAMPS models in the same fashion as the real event.

Performing this analysis yields value for the $\log_{10} \mathcal{B}_U^L$ of 0.37 and 0.79 for the isolated point mass and SIS profiles respectively which are lower than those given for the event. This suggests that it would be difficult to confidently confirm an event with

these parameters, and therefore this test does not rule out either the possibility of a genuine microlensing event or a noise fluctuation in the data. This again highlights the need for additional investigations such as the aforementioned full scale injection campaign to given greater context to the significance of the calculated Bayes factors.

Residual Power Examination

An additional means of scrutinising a microlensing candidate event is to examine the residual power that is left within the data when the maximum-likelihood waveform fit from the non-lensed PE carried out on the event is removed from the data. In the case of a genuine microlensing event, one would expect to see remaining oscillating amounts of power in each of the detectors due to the unaccounted for oscillating behaviour of the amplification factor. In the case of a non-lensed event, absent any other systematic errors, one would expect to see the fluctuations associated with the noise of the detector. This type of analysis is also performed when looking for deviations from general relativity, where one looks at the residual power in the data after the maximum-likelihood general-relativity template—equivalent to the unlensed one—has been subtracted from it [25]. The residual power investigation carried out in Abbott et al. [25] for this event yielded a p -value, corresponding to the probability of obtaining a background event with a residual SNR higher than the event, of 0.97. This suggests the remaining power is within expectations for residual noise.

The residuals from performing this subtraction are shown for each of the detectors, for a subset of the total frequency range, in Fig. 10.6.4. As can be seen, none of the detectors display an obvious coherent oscillation in the residual power that would be expected in the microlensing hypothesis. These residuals are more typical of the noise which may indicate that the event is unlikely to be a microlensing event. Hence, despite the increased favouring of the microlensing hypothesis under the SIS case in terms of raw PE analysis, this work draws the same conclusion as that of the lensing searched conducted by the LVK: GW200208, whilst interesting, is not a genuine microlensing event—though it does highlight the need for more systematic studies on the impact of the noise on microlensing searches in the future.

Millilensing Analysis

The range of masses favoured by the microlensing analysis both in Abbott et al. [32] and within this work would also be within the millilensing regime as described in Sec. 10.2. In the analysis performed here, four millilensing waveform models were

used—three with fixed numbers of millilensing signals (two, three, and four signals respectively), and the fourth being a variable multi-signal waveform allowing any number of signals from 1 to 6.

With each of the millilensing waveform models, we performed parameter estimation of the source and lensing parameters using the Bayesian inference library BILBY [49, 305] with the DYNesty [333] sampler and the IMRPhenomXPHM waveform approximant [295], following the method developed in Liu et al. [230].

The plots resulting from these PE runs are presented in this section and Appendix 10.C. Before commenting on each of the results individually, we note the terminology used commonly for each of the figures. The millilensing parameters are described by a series of effective luminosity distances, d_j^{eff} , time delays with respect to the first image t_{j+1} , and Morse phase n_j for the j^{th} image. The convention for this work is that the images are referred to in time ordering.

Turning to the specific results, we begin with Fig. 10.6.5 detailing the two-signal case. The posterior distribution of the effective luminosity distance of the first signal, d_1^{eff} , displays a clear peak as would be expected of a real signal. The posteriors for the second signal parameters—i.e. the effective luminosity distance and the time delay—both display peak-like features but also have an extended underlying posterior. Without additional evidence this would be insufficient to claim millilensing of the signal and could be explained by the presence of noise, as discussed in the microlensing analysis of the event.

The three-signal analysis results (detailed in Appendix 10.C) are in agreement with the two-signal case, where the effective luminosity distance, d_2^{eff} , and time delay of the second signal, t_2 , display a peak-like feature in the posterior distribution. The corresponding parameters of the third signal, however, do not show significant peaks in their distributions which disfavors the presence of a third millilensing component signal. The four-signal analysis (detailed in Appendix 10.C) also lacks any peaking features in the parameters of either the third or fourth signal—returning their uniform priors and giving additional evidence for the disfavouring of any more than two signals.

Lastly, a multi-signal analysis, making the number of millilensing components signals itself a free parameter, was performed. In this analysis, the number of signals was allowed to range from 1 to 6. The posterior distributions for the millilensing parameters are shown in Appendix 10.C. These posteriors are again consistent with

Table 10.6.2: Comparison of Bayes factors for the evidence against the unlensed hypothesis from the millilensing runs for GW200208.

Model	$\log_{10}(\mathcal{B}_U^{\text{Milli}})$
Two signals	0.86
Three signals	0.92
Four signals	0.96
Multi-signal	1.10

the assertion that there is no favouring for any number of signals above a possible second one. The additional results of attempting to infer the number of signals are shown in Fig. 10.6.6. The discrete posterior here is notably ambiguous disallowing confident constraints on the number of signals here—despite only the posteriors of the second image having any notable features. This serves to underline the fact that the features within the second image posteriors are insufficient to claim a millilensing detection.

In addition to the posterior plots, we also compute the Bayes factors for the millilensing hypothesis against the unlensed one. The values are given in Table 10.6.2. They slightly favor the millilensing case, not significantly enough to truly favor this hypothesis when accounting for the astrophysical information, the prior odds, and the observed posteriors.

It is thus the conclusion of the millilensing analysis that there is insufficient evidence to support millilensing within GW2000208 despite a favouring of the millilensing hypothesis when comparing the Bayes factors of signal versus noise for each of the models.

10.7 Conclusions and Prospects

In this work, we have analyzed candidates found to be interesting by the LIGO-Virgo-KAGRA lensing searches in the full O3 data [32] as if they were genuinely lensed. We considered three main types of lensing: strong lensing, millilensing, and microlensing with the types being defined by the effects they have on observed gravitational waves. Though the events investigated do not display strong evidence of being lensed, the analyses done here demonstrate possible follow-up strategies for future observing runs in order to assess the significance of any lensing candidate event.

Firstly for the GW190412 event—which displayed the greatest support for being a type II image—, we analysed the data with two other waveform models, showing that these do not show as strong a feature. Therefore, the observed support is most likely due to combined noise and waveform systematics. Additionally, we study the event with a microlensing pipeline to see if such lensing could lead to the apparent deviations in the overall phasing, finding no evidence for this hypothesis.

Next considered was the GW191103–GW191105 pair, which was flagged as interesting because of its relatively high coherence ratio and the consistency of the relative amplitudes and time separation with the expectations for the relative magnification and time delay of galaxy lenses. Testing the effect of waveform systematics on the posterior overlap analysis showed that the lensing hypothesis is favoured regardless of waveform choice. We then went on to demonstrate that whilst the event is compatible with galaxy-lens models, inclusion of models in the coherence ratio ultimately does not yield a significant increase in support, as seen by the low coherence ratio. The disfavouring of the lensed hypothesis is further shown by including a more realistic SIE model in our analysis pipeline, still finding a negative log Bayes factor. Furthermore, we demonstrated how an electromagnetic counterpart to the host galaxy could be searched for and showed that no confident counterpart could be found. We also demonstrated that neither of the individual events has any indications of microlensing effects. Finally, we looked for a subthreshold lensed counterpart but found no promising candidate.

A new ranking scheme for the sub-threshold counterparts of detected supra-threshold events found a new interesting candidate pair: the GW191230_180458 supra-threshold and the LGW200104_180425 sub-threshold events. We performed investigations with additional dedicated sub-threshold searches which confirmed interest in the event pair. As was done for the other event pairs, we then analysed it using the standard and follow-up tools. First with the posterior overlap analysis, we saw the event pair is an interesting candidate. Again, a waveform systematic study yielded consistent results for the various waveforms considered in this work. In this case, analysis with the joint parameter estimation tools showed that upon the inclusion of a galaxy-lens model, the coherence ratio was higher than for the previous pair though only to the extent that 40 unlensed events can produce a pairing with similar results by coincidence. Additionally, performing the computation of the proper Bayes factor with an SIS model leads to negative log Bayes factor, disfavouring the lensing hypothesis. On the other hand, the inclusion of an SIE lens model leads to

a marginally positive log Bayes factor. However, it is not high enough to compensate for the prior odds, and therefore the posterior odds is disfavoring the lensing hypothesis. Besides, the low p_{astro} and FAR cast doubt on the astrophysical origin of the sub-threshold event. Finally, as with the previous pair, a search for possible electromagnetic counterparts yielded no confident matches which is in line with the expectation for the events not being lensed. Let us also re-iterate the absence of clear evidence for the sub-threshold event to be genuine in the first place.

The last event analysed was GW200208_130117 which was flagged as the event closest to the expectations for a microlensing event. First, we re-examined the event using different lens models. We found the Bayes factors to be slightly higher than those computed in Abbott et al. [32], but still compatible with values found for background unlensed events. We also studied the variation of the Bayes factor for a point-mass lens model depending on the priors used. We found values ranging from slightly disfavoring lensing to favoring this hypothesis, in line with other analyses performed on this event. To deepen the investigation, we perform a maximum-likelihood injection, recovering the injection with a Bayes factor comparable to the one found for the event, showing the difficulty to confidently identify microlensing at this sensitivity. We then looked at the residual power remaining when subtracting the best-fit unlensed waveform. This test did not yield any particular evidence for the remaining power being consistent between the detectors, which is in line with the event being most likely not microlensed. Finally, since millilensing may also lead to beating patterns in the waveform, we searched for millilensing features in this event. These searches demonstrated that there was no additional evidence for any more than two lensed waveforms comprising the event, and the combination of posteriors and Bayes factors were not sufficient to conclusively favour the millilensing hypothesis in general.

In the end, the conclusions from the additional tests are in line with those given by the LIGO-Virgo-KAGRA Scientific Collaboration [32], showing that none of the events or event pairs is likely to be genuinely lensed, regardless of their initially intriguing characteristics. By doing these additional studies, we have shown some important points for future lensing searches, such as the possibility of having waveform systematics, the impact of the lens model in the analysis, the difficulties one may have to distinguish between events resembling each other by chance and genuinely lensed ones, the interplay between microlensing and millilensing, and other additional avenues to further investigate lensing candidates in the future. These

follow-up methods should be valuable in the future when more intriguing lensed candidates are found.

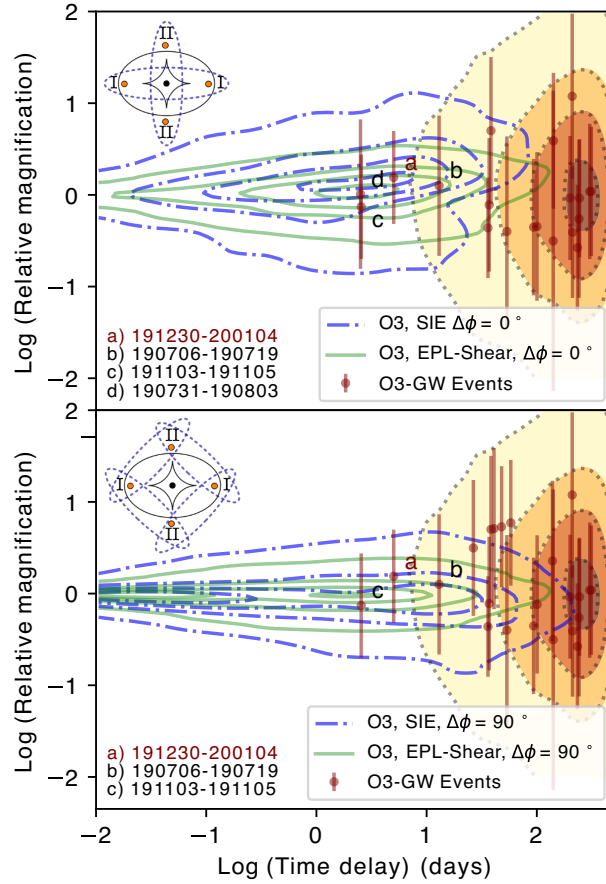


Figure 10.4.3: Comparison of the posterior for the observed relative magnification and time delays for the O3 event pairs with the expected distributions for the lensed population of mergers from More and More [259] (dashed blue, using an SIE model) and from Wierda et al. [364] (solid green, using an elliptical model with a power law density profile and external shear (EPL-Shear)) and unlensed population (yellow-orange-red), both assuming galaxy lenses. Overlaid in brown are the observed values for selected O3 event pairs, and the letters mark the event pairs more compatible with the lensing hypothesis. Written in brown, and denoted with letter a, is a pair made of a super-threshold and a sub-threshold event, and further analysed in this work. The top panel corresponds to the expected distribution when the two images are of the same type, i.e., there is no phase difference between the two images (see top-left illustration), while the bottom panel corresponds to a configuration where the two image types differ, i.e., there is a $\pi/2$ shift between the images. Most of the observed event pairs are well outside the lensed distribution. The GW191103–GW191105, GW190706–GW190719, and GW191230–LGW200104 pairs are more compatible with the lensed hypothesis than with the unlensed one. In particular, the GW191103–GW191105 pair lies in a higher probability density region than the other pairs. One also sees that the GW191230–LGW200104 pair —made of a supra and a sub-threshold event— lies in a higher density region, even if it is less important than the GW191103–GW191105 pair. This pair is discussed further in Sec. 10.5.




Figure 10.4.4: From dark to lighter, the 10%, 50%, and 90% confidence sky localisation for the GW191103–GW191105 pair. Overlaid are the cross-matched 5 candidates from the Master Lens Database.




Figure 10.4.5: Posteriors for a subset of detector-frame source parameters and the lensing parameters produced during the GRAVELAMPS microlensing analysis of GW191103. As can be seen, whilst the source posteriors are well constrained, the lensing parameter posteriors are extremely broad and uninformative. This shows there are no observable microlensing features in this signal.




Figure 10.5.1: Scatter plot of posterior samples of GW191230, marginalised to the mass and aligned spin parameter space. We choose the template given by a KDE estimator. The other two options shown are the maximum likelihood point (ML) and an alternative KDE with an extra set of samples reflected on the $m_1 = m_2$ line (rKDE).




Figure 10.5.2: Overlaid LGW200104 and GW191230 skymaps with 90% and 50% confidence regions.

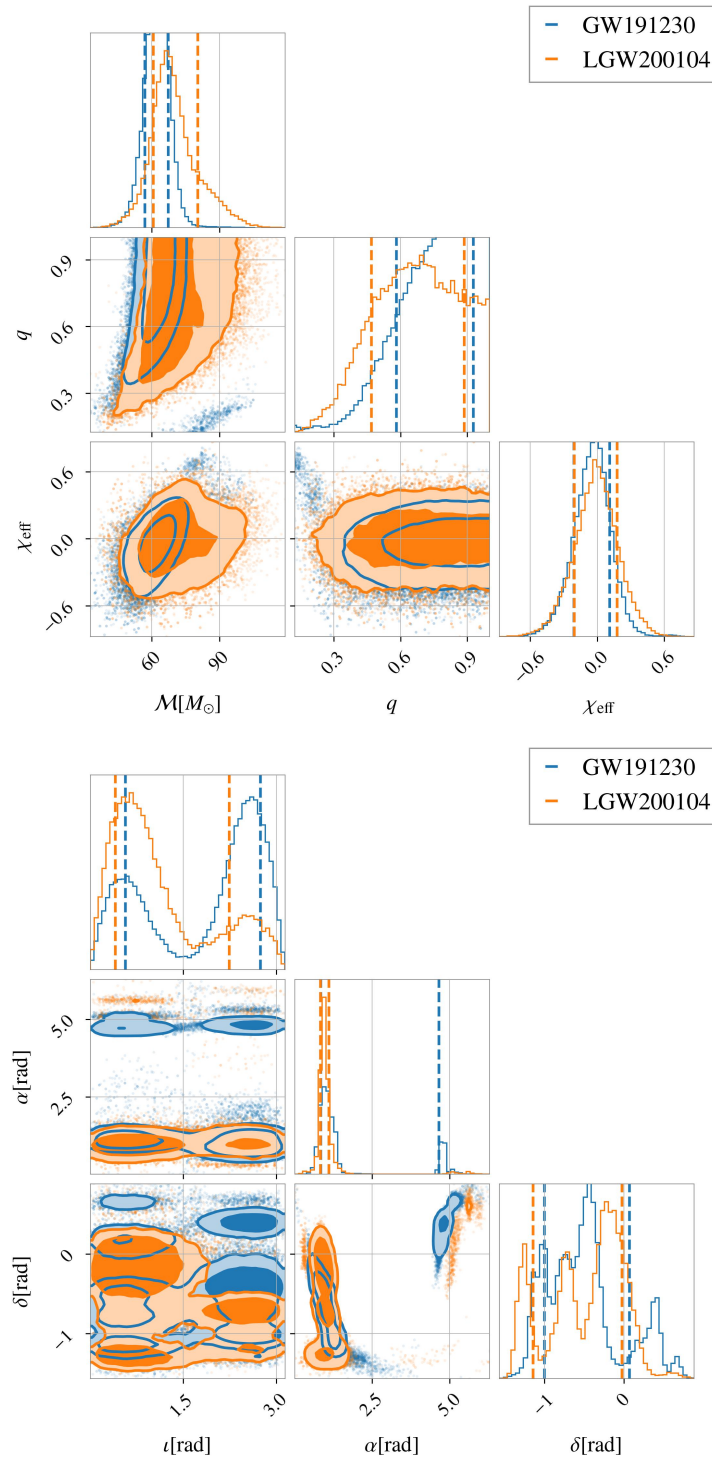


Figure 10.5.3: Posteriors for GW191230 (blue) and LGW200104 (orange). The posteriors, though broad, have significant overlap for both the intrinsic (left) and extrinsic (right) parameters.




Figure 10.5.4: The difference in Morse factor for the GW191230–LGW200104 event pair according to the GOLUM pipeline. The preferred value is 0.5 but there is also some support for 0.




Figure 10.5.5: Posterior distribution for the relative magnification for the GW191230–LGW200104 event pair measured with the GOLUM (solid blue) and HANABI (dashed orange) pipelines. One sees that the measured values are consistent between the two pipelines.

image type of LGW200104_18	III	0.0%	0.0%	16.6%
	II	0.0%	14.3%	26.2%
	I	14.0%	28.9%	0.0%
		I	II	III
		image type of GW191230_18		

Figure 10.5.6: Posterior probability mass function for the image type of GW191230 and the image type of LGW200104 from HANABI. It is consistent with the GOLUM result that it is more likely for the difference in Morse factor to be $\Delta n = 0.5$ (i.e. the II-I and III-II configuration) than to be $\Delta n = 0$ (i.e. the I-I, II-II and III-III configuration).




Figure 10.5.7: Sky localisation 10%, 50%, and 90% confidence region (from dark to light) for the GW191230-LGW200104 pair. Overlaid are the cross-matched 21 candidates from the Master Lens Database.

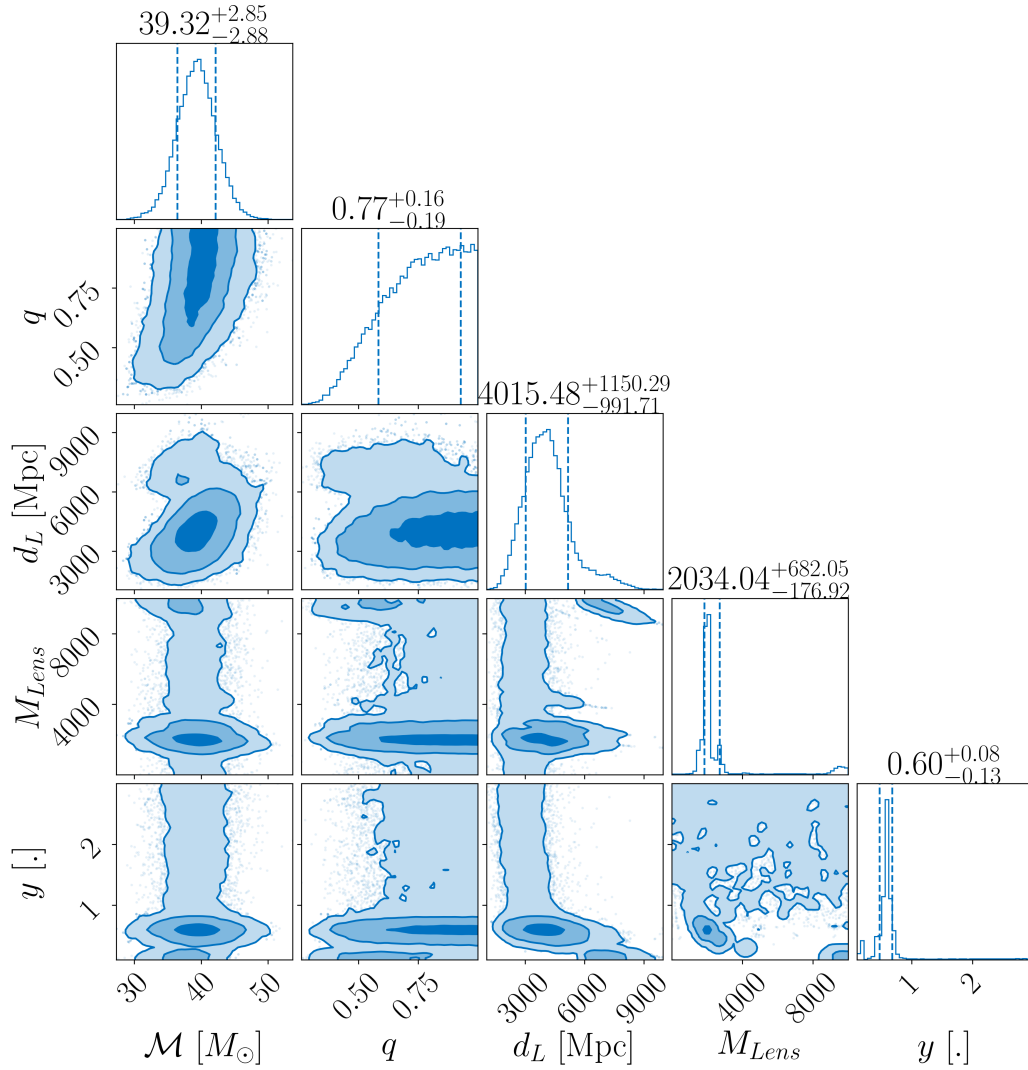


Figure 10.6.1: Posteriors of a subset of source parameters as well as lensing parameters for GW200208 in the SIS microlensing case. Unlike the other events that have been examined within this work, the lensing parameters for the model are well constrained, even if the $3\text{-}\sigma$ confidence interval is a bit noisy. This means that this event, unlike the others, cannot be immediately ruled out as a lensing candidate from this test.

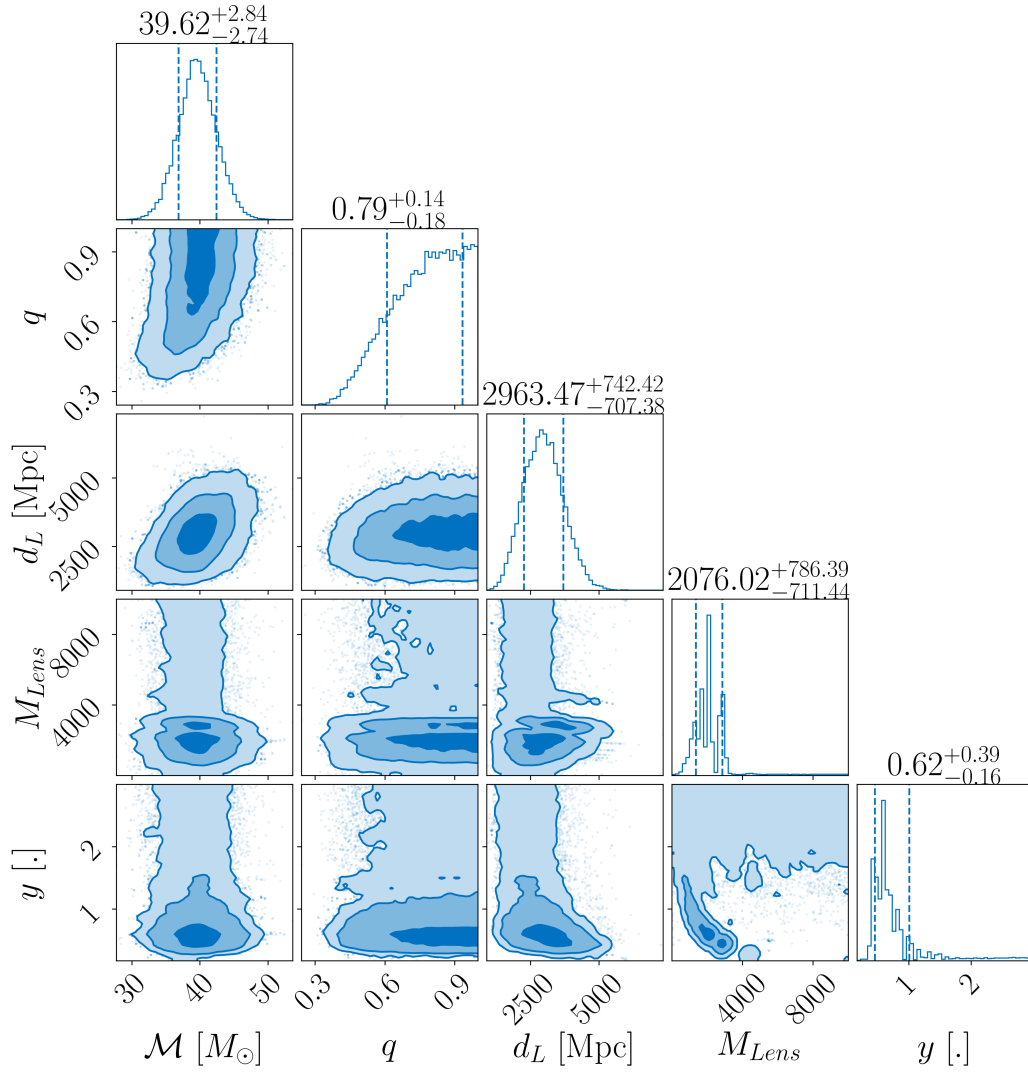


Figure 10.6.2: Posteriors of a subset of source parameters as well as lensing parameters for GW200208 in the point mass microlensing case. For this case, the posteriors are similarly constrained, notably arriving at similar lensing parameters as the SIS case even if the $3\text{-}\sigma$ confidence interval is a bit noisy.




Figure 10.6.3: The posterior densities of the recovered microlensing parameters for different parameter estimation runs, as presented in Table 10.6.1. The results are visualized with varying colors from light to dark (numbered from 1 to 11, indicating different runs as we move down the table.)




Figure 10.6.4: The residual power remaining, from top-to-bottom, in LIGO Hanford, LIGO Livingston, and Virgo when subtracting the best fit non-lensed waveform template for GW200208 as determined by the unlensed PE from the detector strain, over a subset of the total frequency range. As can be seen, there is no obviously coherent oscillatory behaviour in the residual power which would be expected in the case of the microlensing hypothesis. This absence would suggest that what remains is noise related rather than signal related.




Figure 10.6.5: Corner plot of the millilensing parameters obtained from a two-signal analysis of GW200208. Notable is that there is a peak within the effective luminosity distance and time delay parameters for a potential second millilensing signal. However, this peak could be explained by the presence of detector noise.




Figure 10.6.6: Discrete posterior of the number of signals in the multi-signal analysis for GW200208. This posterior is insufficient to confidently assert a number of signals present.

APPENDIX

10.A LENSID GW191103–GW191105 Investigations

In Abbott et al. [32] the pairs which had false positive probability (FPP) less than 0.01 either with PO or LENSID were passed on for the follow-up analysis. According to PO the GW191103–GW191105 pair is found to be one of the most significant candidates. However, the LENSID FPP is found to be 0.16. We cannot determine for certain why LENSID did not find the pair significant for follow-up analysis, however, we can identify some possible contributing factors. Before detailing these, we briefly summarise how LENSID works. LENSID is made up of two ML models, one which takes Q-transforms input, and another which takes skymaps as input. On the basis of the Q-transforms, the network outputs the probability for the event pair to be lensed for each detector. Additionally, there is one output lensing probability based on the sky map. The entire probability for lensing is then computed by taking the four individual probabilities mentioned above. For more details we refer the reader to Goyal et al. [177].

GW191103 was observed only in two detectors, LIGO Hanford (H1) and LIGO Livingston (L1), whereas, GW191105 was observed in all three detectors but was contaminated by a glitch in the Virgo detector. As seen in Fig. 10.A.1, the final PE skymap of the event (right panel in Fig. 10.A.1), which is made after deglitching the data [30], is different from the initial skymap (left panel in Fig. 10.A.1), reducing the sky map FPP from 0.08 to 0.02 after using the PE sky map, still not crossing the threshold.

For the Q-transforms, only the H1 and L1 detectors data are used by the framework. We notice that the Q-transforms for the events, especially for GW191105, are visually poor. They seem to be broken in the middle, as shown in the Fig. 10.A.2. Notice that the Q-transform of GW191105 in the L1 detector has a gap in the middle of the signal with peaks of power on both sides of the gap. This is not expected from a GW chirp signal. We checked that even though the SNRs are similar for both the events in the H1 and L1 detectors, the estimated probability for lensing varies a lot between the two detectors, 0.86 for H1 and 0.12 for L1. This indicates that the ML algorithm is not robust to real noise fluctuations, which is expected as it is trained using simulated Gaussian noise signals. Additionally, from an injection study we found that LENSID is more prone to misclassifying lensed signals with low chirp masses ($< 20M_{\odot}$) which is the case here. In the future, to mitigate these problems,

the ML models will be trained and tested on data containing real noise and lower chirp masses.




Figure 10.A.1: Comparison of sky maps for GW191105, as downloaded from GraceDB which is generated using LALINFERENCE (left) and from GWOSC which is generated using BILBY after de-glitching the Virgo data (right). The LALINFERENCE sky map is narrower as compared to BILBY, likely because of the glitch present in the data. Note that the localisation patch near to the north pole is the one that overlaps with GW191103 and therefore matters more than the rest of the patches.

10.B Discarded Targeted PyCBC Sub-Threshold Search Triggers

In this section, we show the time-frequency maps for the two discarded (third and fourth ranked) candidate triggers found as possible sub-threshold counterparts for the GW191230 event by the PyCBC-based pipeline. The two can clearly be identified as glitches, with the third-ranked clearly having a power excess across a broad frequency band at the same time without presenting a time-frequency evolution similar to the one expected for a genuine GW signal (see Fig. 10.B.1), and the fourth in ranking (see Fig. 10.B.2) clearly matching a scattered-light glitch [327, 328, 343].

10.C Details of the Millilensing Investigation

This Appendix presents additional details of the millilensing investigation for GW200208.

Fig. 10.C.1 represents the result for a millilensing run done with 3 possible superposed images. It shows that in comparison with Fig. 10.6.5 the addition of a third




Figure 10.A.2: Q-transforms images input to the `LENSID` pipeline for GW191103 (top panel) and GW191105 (bottom panel). The chirping feature for GW191105 is broken in both the LIGO detectors, whereas for GW191103 the chirp signal is fairly visible in Hanford, and not so visible in Livingston. This could be one of the reasons why `LENSID` did not identify this pair as significant.




Figure 10.B.1: Time-frequency map of the third-ranked PyCBC candidate. This shows the glitch present near to GPS time 1260223739. This represents the kind of quieter glitches that gets skipped in the normal autogating procedures.




Figure 10.B.2: Time-frequency map of the fourth-ranked PyCBC candidate, consistent with a scattered-light glitch [327, 328, 343].

image is not leading to the recovery of an extra possible image since the posterior for its lensing parameters are flat and uninformative. Similarly, the posteriors for the four-image analysis (Fig. 10.C.2) show flat posteriors for the lensing parameters of the third and fourth possible images, leading to the conclusion that no more than two images can be identified in the data.

In addition to a posterior on the possible number of images, the run where the number of images is left free also returns posteriors for the lensing parameters of the different images. These are shown in Fig. 10.C.3. Only the posteriors for a possible second image are not completely uninformative. The others are flat, meaning that the analysis does not favour anything with more than two signals.




Figure 10.C.1: Corner plot of the millilensing parameters obtained from a three-signal analysis. A similar peak appears in the second image parameters as was present within the two-signal analysis shown in Fig. 10.6.5. However, no such features are present within the posteriors of the third millilensing image.

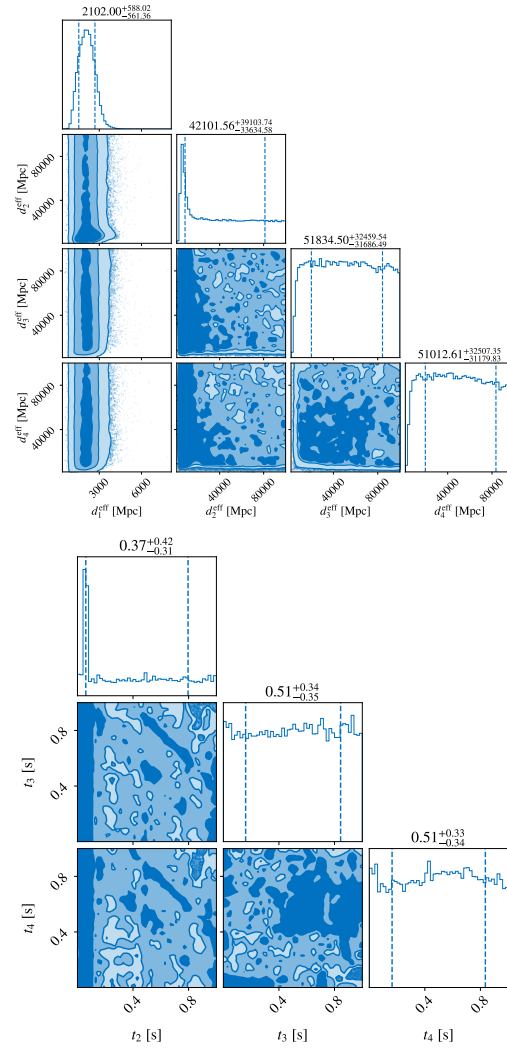


Figure 10.C.2: Corner plots of the effective luminosity distance (left) and time delay (right) parameters obtained from the four-signal analysis. Consistent with the previous analyses, there are no peaking features in the third or fourth signal posteriors.

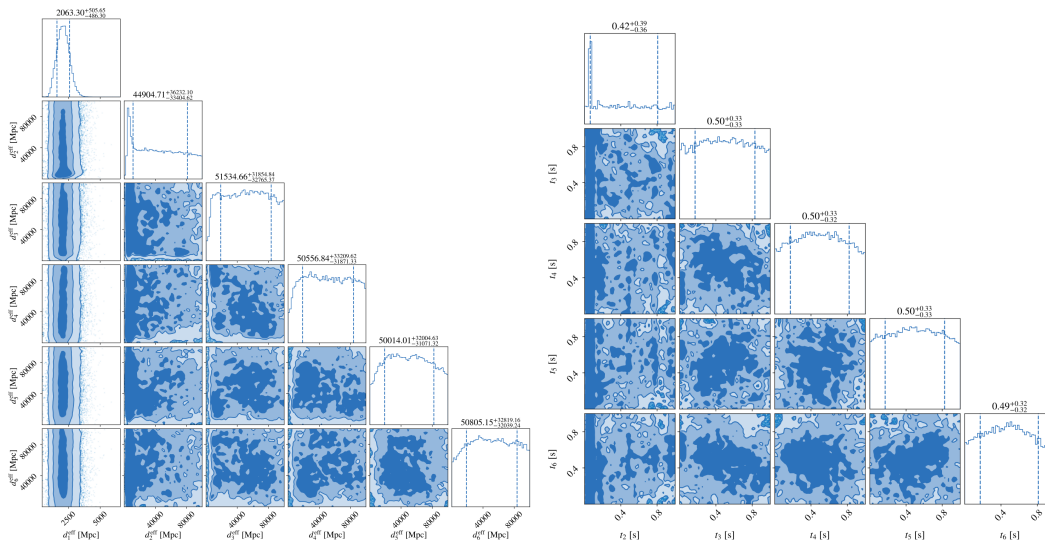


Figure 10.C.3: Corner plots of the effective luminosity distance (left) and relative time delays (right) obtained from a multi-signal analysis. Again, these are consistent with what has been seen in the previous analyses with no favouring for any number of signals above two.

Chapter 11

TESLA-X: A MORE EFFECTIVE TARGETED SEARCH METHOD FOR SUB-THRESHOLD STRONGLY LENSED GRAVITATIONAL WAVES

Note: This Chapter is an adaption of the short author preprint:

“*Alvin K. Y. Li et al., “TESLA-X: An effective method to search for sub-threshold lensed gravitational waves with a targeted population model”, arXiv e-prints, 2023, DOI: [arXiv.2311.06416](https://arxiv.org/abs/2311.06416).*”

Alvin K. Y. Li is the lead author of this paper, which describes the next-generation “TargetEd Subthreshold Lensing seArch” (TESLA) method with the use of a better-constructed reduced template bank and a targeted population model.

11.1 Introduction

Strongly lensed gravitational waves will likely originate from sources living in the higher redshift universe. Without being lensed, these signals will only be barely detectable or undetectable. Under strong lensing, copies of gravitational waves from these high redshift sources will be produced; some may be magnified and are detected as super-threshold gravitational waves, but some may be potentially demagnified with a reduced signal-to-noise (SNR) ratio below the detection threshold. These signals, formally known as “sub-threshold” signals, are of particular interest, as a significant proportion of strongly-lensed gravitational-wave events falls into this category [364]. Retrieving these possible sub-threshold lensed counterparts is a crucial step in boosting confidence for the first detection of strongly-lensed gravitational waves.

Two targeted search methods [223, 247] have been developed and employed to recover potential sub-threshold lensed counterparts to known super-threshold gravitational waves. The GstLAL-based TargetEd Subthreshold Lensing seArch (TESLA) described in [223] strategically reduces the search parameter space to look for possible sub-threshold lensed counterparts. However, as we will explain in this chapter, the traditional TESLA method can potentially lose signals due to an excessive loss in signal-to-noise (SNR) ratio due to how the targeted template bank is constructed.

Additionally, TESLA's efficiency in finding potential sub-threshold lensed gravitational waves can be improved if a targeted population model built based on information from the target superthreshold event is used. A proper targeted population model in the ranking statistic can further reduce the influence of the noise background. Detailed information on the TESLA method can be found in [223].

In this chapter, we present a novel method, TESLA-X, to search for strongly-lensed sub-threshold gravitational-wave signals. TESLA-X employs a targeted population model to enhance the search sensitivity for potential lensed sub-threshold signals, and it addresses the excessive SNR loss issue by utilizing a new targeted template bank construction method.

This chapter is organized as follows: Section 11.2 outlines the workflow of the traditional TESLA method. In Section 11.3, we give an overview of how the population prior term in the likelihood ratio calculations for triggers from the GstLAL search pipeline is calculated, and detail the TESLA-X methodology. Section 11.4 presents a simulation campaign to evaluate the TESLA-X method. Finally, in Section 11.5, we summarize our findings and discuss potential future work to further enhance the search sensitivity of the TESLA-X pipeline.

11.2 The search for sub-threshold strongly-lensed gravitational waves: Objectives and current approach

Current effort to search for strongly-lensed sub-threshold lensed counterparts to super-threshold gravitational waves

When searching for gravitational waves in general, signals with sufficiently high amplitudes that can be identified as gravitational waves from the noisy data are known as *super-threshold*. In contrast, weaker gravitational waves not distinguishable from noise are known as *sub-threshold*. As discussed in Chapter 6 and 7, strong lensing can produce multiple copies of gravitational waves coming from the same source with different amplification factors $\sqrt{\mu_j}$. $\sqrt{\mu_j}$ can take on values smaller than 1, i.e. the repeated gravitational-wave signals are *de-magnified* compared to the not-lensed gravitational-wave signal and may become sub-threshold. It is then natural to ask the question: If we assume that one of the identified super-threshold gravitational waves detected so far is strongly lensed, can we find its sub-threshold strongly lensed counterparts, if they exist?

There have been ongoing efforts to employ *matched-filtering based search pipelines*

with modifications [223, 247] to search for strongly-lensed sub-threshold lensed counterparts to identified super-threshold gravitational waves from LVK’s first three observing runs (O1, O2 and O3) [23, 32, 186]. In this paper, we focus on the *TargetEd subthreshold Lensing seArch (TESLA)* method [223], which is built based on the GstLAL search pipeline [43, 86, 96, 176, 183, 249, 250, 263, 309]. The following sub-section gives a brief introduction to the traditional TESLA method.

The traditional TESLA search method

Traditional TESLA is built based on the matched-filtering-based pipeline GstLAL. A detailed introduction to the GstLAL search pipeline can be found in [86, 223, 250]. In a general search for gravitational waves from compact binary coalescences, the goal is to detect signals that span the entire parameter space to which the detector network’s frequency bandwidth is sensitive. Hence, a large template bank comprising over a million template waveforms is constructed and used. A large number of candidates (also known as triggers) is generated, regardless of whether they are noise or signals. The more templates we use in a search, the larger the trials factor, and hence, the larger the noise background will become.

Gravitational waves with weaker amplitudes, for instance, strongly-lensed sub-threshold lensed counterparts, can be buried in the large noise background. Hence, it is necessary to tactically reduce the noise background while keeping the targeted foreground constant to uncover possible weaker gravitational waves.

Assuming a super-threshold gravitational wave is strongly lensed, we want to search for its possible sub-threshold lensed counterparts. In Chapter 7 we explained that strongly-lensed gravitational waves from the same source should have identical intrinsic parameters (e.g. component masses and spins), and hence identical waveforms. For each super-threshold event (target), we obtain the posterior probability distribution that gives the best estimates of the source parameters of the target using the Bayesian parameter estimation analyses described in [48, 49, 305, 356]. We can then narrow down the parameter search space for possible strongly-lensed sub-threshold counterparts to regions consistent with the target’s posterior parameter space. However, as explained in [223], the posterior parameter space of the target itself is insufficient because of noise fluctuations in the data. Non-Gaussianity in the data can cause false signals (especially those that have weaker amplitudes and are hence sub-threshold), to be registered by templates that have very different parameters than those of the posterior samples for the target.

TESLA accounts for both the signal sub-space and noise fluctuations in the data when searching for possible strongly-lensed sub-threshold counterparts to super-threshold gravitational waves. Here, we outline the steps for running a TESLA analysis. For each super-threshold gravitational event (target), we start with the posterior samples ordered in decreasing order of log-likelihood from the parameter estimation for the target. Since the amplitudes of possible strongly-lensed sub-threshold gravitational waves are lower, they will be registered as triggers in the search pipeline with lower signal-to-noise ratio ρ (SNR). For each posterior sample, we generate 1 injection with the exact same parameters (including luminosity distance to the source D_L and sky location, i.e. right ascension α and declination δ) as the sample, and 9 extra injections with D_L increased (and hence optimal SNR ρ_{opt} decreased) to mimic the de-magnifying effect of strong lensing on the amplitudes of possible sub-threshold lensed counterparts, according to the inverse proportionality of D_L and the optimal SNR ρ_{opt} :

$$D_L \propto \frac{1}{\rho_{\text{opt}}}. \quad (11.1)$$

We require the latter 9 injections to have detector SNR $\rho_{\text{det}} \geq 4$ in at least one detector as a constraint from the GstLAL search pipeline ¹. The set of lensed injections represents the possible sub-threshold lensed counterparts to the target we are searching for. They are then injected into real data, and we perform an injection campaign using GstLAL with a full template bank to recover them. After the injection campaign, templates that can find these injections ² are kept to construct a targeted template bank to search for possible strongly-lensed sub-threshold counterparts to the target. The end product of TESLA is a ranked list of possible lensed counterparts to the target in increasing order of FARs. Readers are reminded that the FARs (or other ranking statistics) reported in the search's ranked list do not indicate how likely the signal candidate is a lensed counterpart to the target, but only a priority label for follow-up analysis.

For a detailed explanation of the TESLA method, please refer to [223].

¹ We do not lower the single detector SNR threshold in the GstLAL search pipeline since this will lead to an exponential increase in the number of both signal and noise triggers. As future work, we will explore the possibility of further lowering the single detector SNR threshold.

² An injection is “found” if the FAR of the associated trigger is ≤ 1 in 30 days.

11.3 The upgraded TESLA-X search method

The missing pieces in the traditional TESLA method: Non-optimal template bank and population model

We have shown in [223] that the traditional TESLA method succeeds in improving the search sensitivity to possible strongly-lensed sub-threshold gravitational waves. However, the whole picture still needs to include two puzzle pieces.

First, in the original full template bank, the parameter ranges are informed by the masses and spins of detectable sources (BNS, NSBH, BBH), and the templates are distributed with a density that allows a maximum loss of 3% SNR when a gravitational-wave signal is recovered (compared to the optimal SNR³) [263]. However, when a reduced template bank is generated using traditional TESLA, we discard templates that did not recover lensed injections for the target super-threshold gravitational waves. This breaks the original optimality of the full template bank, caused by the finite number of lensed injections that are used to generate the reduced template bank, and strongly-lensed sub-threshold signals whose optimal detector SNR is near threshold (i.e. around $\rho_{\text{opt}} \approx 4.12$) will not be detected. Therefore, it is vital that the optimality of the reduced template bank must be conserved to maximize our search sensitivity towards possible strongly-lensed sub-threshold gravitational waves.

Second, a population model was not set in traditional TESLA. In the analysis, the absence of a population model carries an implicit assumption that every template in the template bank has an equal probability of recovering a possible sub-threshold counterpart. Since our goal is to detect sub-threshold counterparts whose intrinsic parameters are identical to a super-threshold signal, implementing a population model centered around the signal's recovered parameters will increase the pipeline's sensitivity in detecting signals in that region. The population model should also be dependent on SNR to account for sub-threshold signals that are coloured by noise and, thus, could be recovered by a template that is different from that recovered the super-threshold signal.

³ The optimal SNR is the SNR when the signal is cross-correlated with a template with the exact same parameters as the signal.

The likelihood ratio in the GstLAL search pipeline

To rank candidates, GstLAL assigns each trigger a likelihood ratio as its ranking statistic, defined as [86, 250, 309]:

$$\mathcal{L} = \frac{P(\vec{D}_H, \vec{O}, \vec{\rho}, \vec{\xi}^2, [\Delta\vec{t}, \Delta\vec{\phi}] | \vec{\theta}, \text{signal})}{P(\vec{D}_H, \vec{O}, \vec{\rho}, \vec{\xi}^2, [\Delta\vec{t}, \Delta\vec{\phi}] | \vec{\theta}, \text{noise})} \cdot \mathcal{L}(\theta), \quad (11.2)$$

where the first term, $P(\dots | \vec{\theta}, \text{signal})/P(\dots | \vec{\theta}, \text{noise})$, is the ratio of the probability of obtaining the candidate event under the signal model to that under the noise model. It depends on (1) the detectors \vec{O} that are participating at the time of the candidate event, (2) the horizon distances of the participating detectors \vec{D}_H , (3) the detector SNRs $\vec{\rho}$ and (4) The candidate event's auto-correlation-based signal consistency test values at each detector $\vec{\xi}^2$. If a candidate event is found in coincidence at multiple detectors (i.e. the event is a *coincident* event), \mathcal{L} will also include (5) the relative time delays $\Delta\vec{t}$ and (6) the relative phase delays $\Delta\vec{\phi}$ of the candidate event between participating detectors ⁴.

The second term in Equation 11.2,

$$\mathcal{L}(\theta) = \frac{P(\vec{\theta} | \text{signal})}{P(\vec{\theta} | \text{noise})}, \quad (11.3)$$

is the ratio of the probability that a trigger is recovered by a template with parameters $\vec{\theta}$, given the trigger is an astrophysical signal or noise. Because this ratio does not depend on the data observed by the detectors and depends on input parameters given before performing the search, it has been written as a separate term.

The method of calculating the numerator and denominator are described in detail in [166] and [85], respectively. While the derivations are not repeated here, we wish to highlight that information about the population model is enfolded into the numerator; therefore, it is $P(\vec{\theta} | \text{signal})$ that allows the pipeline to perform targeted source searches. The following section describes the changes made to improve the strongly-lensed gravitational-wave search.

A better reduced template bank, and a targeted population model

Traditional TESLA aims to reduce the noise background effectively to uncover possible strongly-lensed counterparts to the target super-threshold gravitational wave.

⁴ These relative time delays and phase delays should not be confused with the relative time delays and phase delays coming from strong lensing

We reduce the number of templates, and hence trials factors, by keeping only templates that can cover the parameter space where possible lensed counterparts may live in. Through an injection campaign, traditional TESLA locates a subset of templates from the full template bank to construct a reduced template bank to search for possible lensed counterparts to a given target super-threshold gravitational wave. However, as explained at the beginning of 11.3, there are two problems in traditional TESLA, both of which lead to a decrease in the search sensitivity of TESLA.

We now propose two improvements to the TESLA method that will resolve these problems: (1) constructing a better reduced template bank, and (2) generating a targeted population model.

We recall that the sole purpose of the injection campaign is to determine the search parameter subspace that we should focus on to look for possible strongly-lensed sub-threshold counterparts to the target super-threshold gravitational waves. The injection campaign is always imperfect because we can only do a limited number of injections with discretized SNRs, i.e. we cannot cover the full spectrum of SNRs that possible lensed sub-threshold signals can have. Therefore, the templates that can recover lensed injections in the injection campaign can only serve as a pointer as to the region of search parameter space we should be targeting, i.e. a search parameter subspace where possible lensed sub-threshold counterparts to the target super-threshold gravitational waves may live in. We should hence use the injection campaign results to *define a boundary* within which templates live should be used collectively to construct the reduced template bank to search for possible lensed counterparts to the target, regardless of whether the templates can recover injections from the injection campaign. The proposed reduced template bank does not discard templates within the narrowed-down search parameter space defined by the boundary. Hence, it maintains the optimality guaranteed by the original full template bank within the limited search space. In turn, any possible lensed counterparts that live within the boundary will be recovered with an SNR loss less than 3%, preventing any loss of potential signals due to excessive SNR loss.

Under the strong lensing hypothesis, possible lensed counterparts should have the same parameters as the target super-threshold gravitational waves regardless of being super-threshold or sub-threshold. Therefore, the population of possible lensed counterparts to the target gravitational waves should have a higher probability of lying near the signal subspace defined by the posterior parameter probability distribution of the super-threshold event, and a lower probability of living far away from

that. When we perform the injection campaign in traditional TESLA, we can obtain both the recovered injections' SNRs and the corresponding templates that recover the injections. Generally speaking, sub-threshold signals with lower optimal SNRs are affected by noise fluctuations more substantially and hence can be recovered by templates with parameters different from those of the target super-threshold event. That said, the templates closer to the signal subspace defined by the posterior parameter probability distribution of the target are associated with recovered injections with higher SNR values. Therefore, we can make use of the results from the injection campaign to construct a *Gaussian Kernel Density Estimation* (KDE) function $KDE(\vec{\gamma})$ using the distribution of the templates that recovered the lensed injections, with the corresponding injections' recovered SNRs as weights. The KDE is a function of source parameters $\vec{\gamma}$, which, to match the parameters of the template bank, are set to be the masses and spins of the source. When weighted with the recovered SNRs, regions closer to the signal subspace that the target super-threshold gravitational wave lies in will give a higher value for the KDE function, which matches our expectation that lensed counterparts to the target should have similar parameters as the target. The KDE function can then be redefined as the probability density function that describes the “source population” for our possible strongly-lensed sub-threshold counterparts to the target gravitational wave. Using this as our population model, we then follow the steps detailed in [166] to calculate $P(\vec{\theta}|\text{signal})$, i.e. the probability that a signal is recovered by a template with parameters $\vec{\theta}$. Notably, this probability depends on SNR, such that the search does not overly penalize templates far away from the signal subspace for weak sub-threshold signals. These methods instill a correct population model to the new TESLA method compared to the incorrect uniform population model in traditional TESLA, improving the search sensitivity.

The flow of the TESLA-X search method

We now describe in detail the actual flow of the proposed TESLA-X search method.

Figure 11.3.1 presents the workflow of the TESLA-X method. We start with a target super-threshold gravitational-wave event found during a general search with GstLAL using a full template bank. Using the posterior samples obtained by Bayesian parameter estimation of the target event, we generate, for each sample, one injection that has the same optimal SNR as the sample, and nine additional injections with smaller optimal SNRs by increasing their effective distances, requiring that their SNRs in at least one detector have to be ≥ 4 to ensure they can be registered as

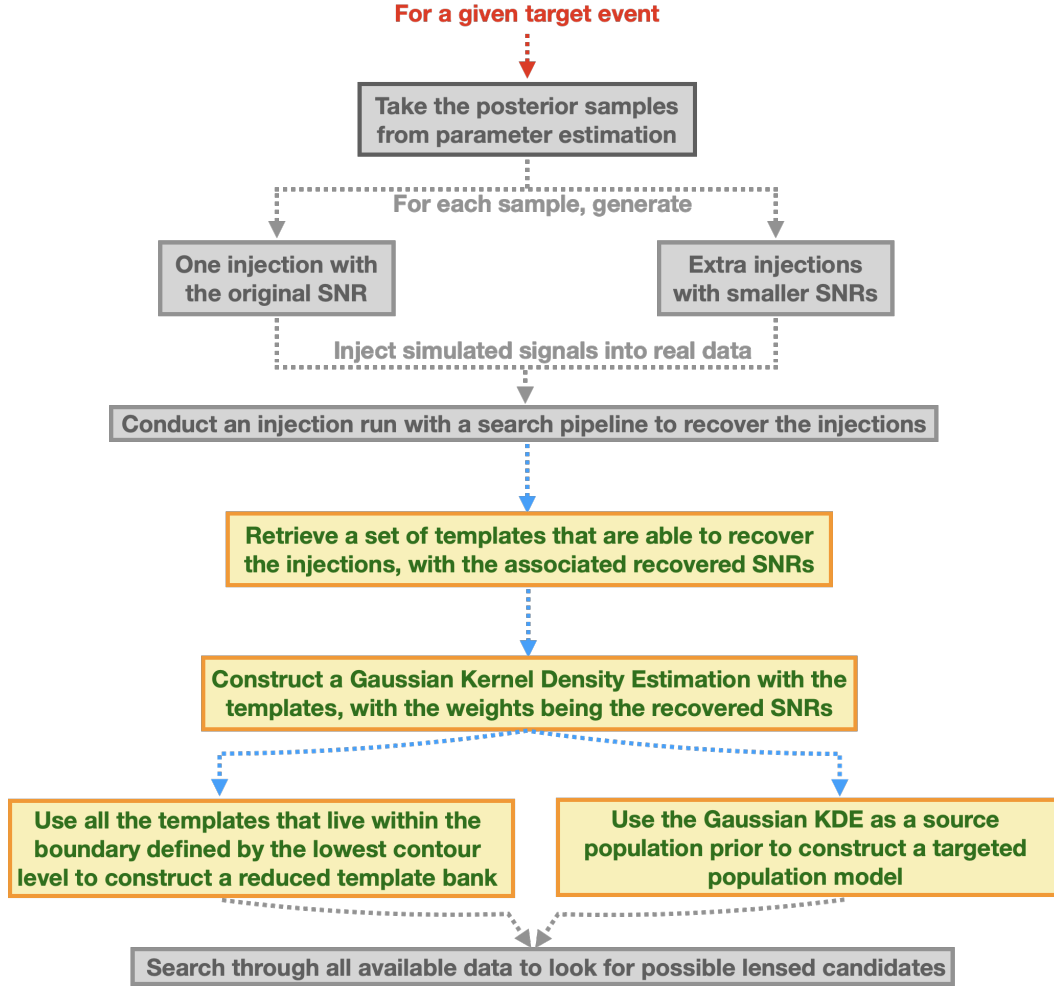


Figure 11.3.1: Workflow of TESLA-X. Elements in grey are the same procedures as in traditional TESLA. Colored elements (color online) are new procedures in TESLA-X.

a trigger during the matched-filtering process in the search. These injections will represent possible sub-threshold lensed counterparts to the target super-threshold event. We then inject these simulated lensed injections into real data, and use the GstLAL search pipeline to try and recover them with a full template bank. Up to this stage, TESLA-X is still identical to the traditional TESLA method.

After the injection campaign, we now retrieve a set of templates that can recover the simulated lensed injections, and the associated observed SNRs of the recovered injections. We can plot the retrieved templates on the chirp mass \mathcal{M}_c - effective spin χ_{eff} space, with the colors of the points being the associated recovered injections' observed SNRs ρ_{obs} . The top left panel of Figure 11.3.2 shows an example cartoon that plots the templates on the \mathcal{M}_c - χ_{eff} space, with the color being the recovered

injections' network SNRs. Traditional TESLA keeps only templates that can recover

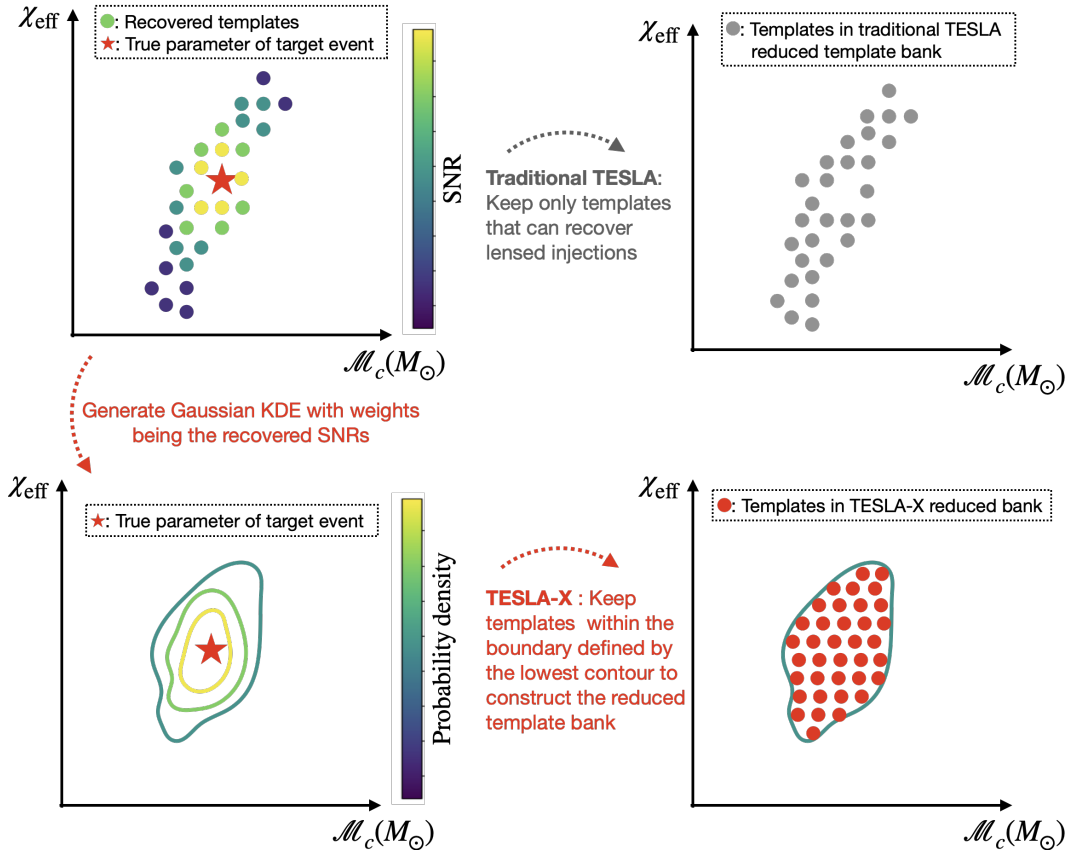


Figure 11.3.2: A cartoon depicting how we construct a reduced template bank from the injection results following the Traditional TESLA method (top right) and the proposed TESLA-X method (bottom left and right).

the lensed injections and use them to construct the reduced bank, hereby known as “TESLA BANK” (top right panel of Figure 11.3.2). For TESLA-X, we construct a Gaussian Kernel Density Estimation (KDE) function $f(\vec{\gamma})$ for the templates that recover the injections, with the corresponding recovered SNRs as the weights, where γ is the set of template parameters we are considering; for example, in Figure 11.3.2 $\vec{\gamma}$ will be chirp mass \mathcal{M}_c and effective spin χ_{eff} ⁵. Using the KDE function, we can then evaluate the probability density for each template point $\vec{\theta}$ in the full

⁵ Note that the lensed injections are generated based on the posterior samples from the parameter estimation of the target event. Spins, admittedly, are not well measured for gravitational waves. However, recall that the lensed injections' spin distribution follows the distribution of the parameter estimation result for the target event. Suppose the spin measurement for the target event is poor. In that case, the parameter estimation result will essentially return the uninformed, uniform prior for spins, meaning that the lensed injections' spins will follow a uniform distribution. Hence, the injection campaign result is not biased towards spins. Hence, we argue that there is no harm in using effective spin as a parameter when generating the Gaussian KDE.

template bank. We then use “matplotlib.pyplot.contour” with the default settings of “matplotlib.ticker.MaxNLocator” to determine the characteristic contour levels for the KDE function (see bottom left panel of Figure 11.3.2) for example). Then, using the lowest contour (it is left as future work to better define the lowest contour / least significant contour in terms of confidence interval) as the boundary, TESLA-X keeps *all* the templates within the boundary and uses them to construct the reduced template bank, hereby known as “TESLA-X BANK”. While the TESLA BANK has “holes” (see, for example, the top right picture in Figure 11.3.2) compared to the original full template bank because we are overly removing templates, the TESLA-X BANK keeps all the templates within a certain boundary, hence recovering the optimality of the original template bank in the region within the boundary. Hence, we will not lose signals due to excessive loss in SNR, and we will expect the TESLA-X BANK to give better search sensitivity.

Because the Gaussian KDE function $f(\vec{\gamma})$ is essentially a probability density function (PDF) that describes the expected distribution of possible (sub-threshold) lensed counterparts to the target super-threshold gravitational wave, we can use it as the astrophysical source population prior to construct a targeted population model using the same methodology described in the previous section and [166]. Figure 11.3.3 shows an example cartoon to illustrate the idea. We have also compared the perfor-

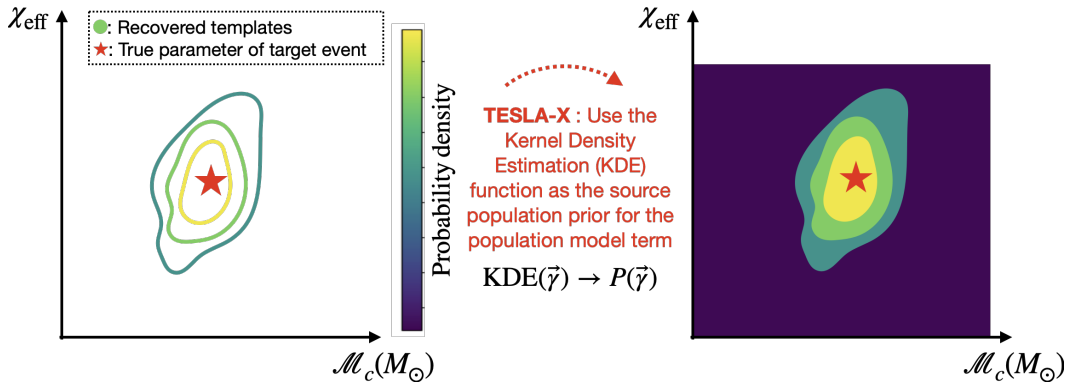


Figure 11.3.3: A cartoon depicting how we use the injection campaign results to construct a targeted population model in TESLA-X.

mance of TESLA-X using (1) the targeted population model suggested above, and (2) an empirical, uniform mass model. We have verified that the search with the suggested targeted population model gives a better performance than the uniform mass model.

Once the reduced bank and the targeted population model have been constructed, we

again use GstLAL to search through all possible data with the targeted bank and the targeted population model to look for potential sub-threshold lensed counterparts to the target super-threshold event. GstLAL will then output a list of candidate events ranked by their assigned ranking statistics, including the FARs and the likelihood ratios \mathcal{L} . We have to stress again that the FARs assigned to the candidate events here do not quantify how likely they are to be lensed counterparts to the target event. Still, they are measures to distinguish noise triggers (false alarms) from real astrophysical signals, whether or not they are lensed counterparts to a target event. The ranking statistics here should only be used as a priority ranking for follow-up analysis to decide how likely each candidate event is a lensed counterpart to the target event. As future work, we will explore the idea of evaluating an additional lensing likelihood ratio for each trigger that includes additional lensing information as a separate ranking statistic. At that stage, the lensing likelihood may reveal how likely a trigger is a lensed counterpart to the target event.

11.4 Method verification and results

In [223], we carry out a full-scale simulation campaign, where we prepare mock data with a pair of super-threshold and sub-threshold lensed signals injected, and apply the TESLA method to recover the two signals. We have proved that the TESLA method can recover sub-threshold lensed signals if they exist. TESLA-X and TESLA have essentially the same working principle, but the only differences are how we construct the reduced template bank and the population model. Therefore, we conduct a simulation campaign solely to compare the search effectiveness and sensitivity towards possible sub-threshold lensed gravitational-wave signals using TESLA and the proposed TESLA-X method.

Mock data generation

First, we generate a 9.15-day long data stream with Gaussian noise recolored with O4 characteristic power spectral densities for LIGO Hanford, LIGO Livingston, and Virgo detector ⁶, with the assumptions that there is no detector downtime ⁷, and no times are vetoed. We inject a super-threshold gravitational wave generated using the IMRPhenomXPHMpseudoFourPN waveform [295] into the mock data. Details about the source parameters of the injected gravitational wave are listed in Table 11.4.1.

⁶ The exact data for O4 characteristic power spectral densities for LIGO Hanford, LIGO Livingston and Virgo detector can be found at <https://dcc.ligo.org/T2200043-v3/public>.

⁷ A detector is considered “down” if it is not in observing mode.

Table 11.4.1: Information of the injected super-threshold gravitational-wave signal MS220508a in the simulation campaign. All properties reported here are measured in the detector frame.

Properties	Injected super-threshold signal
UTC time	May 08 2022 11 : 06 : 00
GPS time	1336043178.397
Distance (Mpc)	2858.18
Primary mass m_1^{det}	$70.08 M_\odot$
Secondary mass m_2^{det}	$38.83 M_\odot$
Dimensionless spins	$\chi_{1x} = 0.182, \chi_{1y} = 0.182, \chi_{1z} = -0.0363,$ $\chi_{2x} = -0.113, \chi_{2y} = 0.132, \chi_{2z} = 0.116,$
Right ascension α	2.811
Declination δ	0.819
Inclination ι	2.513
Polarization Ψ	1.187
Waveform	IMRPhenomXPHMpseudoFourPN

In later parts of this paper, we may refer to the super-threshold signal as MS220508a.

Performing a general search

We then use GstLAL to perform a search over the mock data stream following the same settings used to search for gravitational waves within O3 data in GWTC-3 [27]. We use the same general template bank as described in [223] for the general search. As expected, the search recovers MS220508a with the highest ranking statistics ($\text{FAR} = 2.972 \times 10^{-35}$ Hz, rank 1) among all other triggers. We then apply a Bayesian inference library for gravitational-wave astronomy Bilby [49] ⁸ to perform parameter estimation (PE) for MS220508a, which outputs a set of posterior samples required to apply the TESLA and TESLA-X search pipeline.

Performing an injection campaign

Using the posterior samples from the PE for MS220508a, we generated 7815 simulated lensed injections following equation 11.1. These simulated signals are injected into the mock data, and we perform another search using GstLAL with the general template bank to recover them. The results of the injection campaign are then used to perform the TESLA and TESLA-X analyses.

In the injection campaign, the general template bank recovers 6151 out of 7815

⁸ For reference, we used the default MCMC approach to do the sampling.

injections (See the column “General” in Table 11.4.2), ringing up a total of 1008 templates ⁹. Figure 11.4.1 plots the rung-up templates on the chirp mass \mathcal{M}_c - effective spin χ_{eff} space, with the colors of the markers being the network SNRs of the associated injections. We can see that the results match our expectations: Templates closer to the actual parameters of MS220508a are associated with injections with higher recovered SNRs (i.e. super-threshold). Those that have parameters much different than the actual parameters of MS220508a are, in general, related to injections with lower SNRs (i.e. sub-threshold).




Figure 11.4.1: Rung-up templates by the simulated lensed injections of MS220508a in the injection campaign, plotted on the chirp mass \mathcal{M}_c - effective spin χ_{eff} space. The colors of the markers represent the network SNRs of the associated injections.

Generating the TESLA and TESLA-X reduced template banks

We keep all templates that are rung up during the injection campaign, and use them to generate a TESLA targeted template bank. Figure 11.4.2 shows the templates in the full template bank (in grey) and those in the TESLA template bank (in orange).

⁹ Note that the same template may be rung up multiple times by different injections, and hence the number of rung-up templates can be fewer than the number of found injections.




Figure 11.4.2: The templates in the full template bank (in grey) and in the TESLA reduced template bank (in orange), plotted in the $\mathcal{M}_c - \chi_{\text{eff}}$ space.

For TESLA-X, we first use the rung-up templates in the injection campaign to construct a Gaussian Kernel Density Estimation (KDE) function $\text{KDE}(\vec{\gamma})$ with the associated injections' network SNRs as the weights. γ represents the template parameters used to label the templates. Here, $\vec{\gamma} = \{\mathcal{M}_c, \chi_{\text{eff}}\}$. Then, we evaluate the probability density at each template point in the full template bank, and plot the results on a contour map. Figure 11.4.3 shows the contour map for the Gaussian KDE function estimated for MS220508a.

We then take the lowest contour level (For the case of MS220508a, the lowest contour level has a value of 0.005) as a boundary, and use all the templates within this boundary to construct a TESLA-X targeted template bank. Figure 11.4.4 shows the templates in the full template bank (in grey) and in the TESLA-X reduced template bank (in red) in the $\mathcal{M}_c - \chi_{\text{eff}}$ space. The yellow curve marks the lowest contour level we obtained for the estimated Gaussian KDE function we showed in Figure 11.4.3. There are 22136 templates in the TESLA-X reduced template bank.




Figure 11.4.3: The contour map of the Gaussian Kernel Density Estimation function obtained for MS220508a in the $\mathcal{M}_c - \chi_{\text{eff}}$ space. The colors represent the probability density estimated by the KDE function.

Constructing the targeted population model for the TESLA-X reduced template bank

Next, we use the estimated KDE function for MS220508a and the TESLA-X reduced template bank to construct a targeted population model, following the same steps detailed in 11.3 and [166]. Note that for practical reasons, the Gaussian KDE function $\text{KDE}(\vec{\gamma})$ is re-calculated with $\vec{\gamma}$ chosen to be the component masses and spins m_1, m_2, χ_1, χ_2 instead of \mathcal{M}_c and χ_{eff} .

Performing a re-filtering to recover the lensed injections with the TESLA and TESLA-X methods

Finally, we perform two searches, one using the TESLA reduced template bank, and one using the TESLA-X reduced template bank, together with the targeted population model, over the same stretch of mock data to try to recover the same set of lensed injections used in the injection campaign to compare the performance of the TESLA and TESLA-X methods. Table 11.4.2 summarizes the findings. As expected, the TESLA method successfully recovers more injections than the full




Figure 11.4.4: The templates in the full template bank (in grey) and in the TESLA-X reduced template bank (in red), plotted in the \mathcal{M}_c - χ_{eff} space. The TESLA-X template bank is constructed by keeping templates that fall within the lowest contour level from the Gaussian KDE function shown in Figure 11.4.3, i.e. 0.005, plotted as a yellow curve in this figure. For easy comparison, we plot the TESLA reduced template bank templates in the same figure (in blue).

Table 11.4.2: Number of injections found during the search of mock data using the general template bank, TESLA method, and TESLA-X method, respectively.

Injections	General	TESLA	TESLA-X
Total	7815	7815	7815
Found	6151	6169	6297
Found % change	-	+0.29%	+2.37%

template bank, with an increase of 0.29%. However, we can see that the proposed TESLA-X method outperforms the TESLA method. In particular, it recovers even more lensed injections than the TESLA method, with an increase of 2.37%.

Furthermore, if we look at the sensitive range v.s. FAR for MS220508a-alike signals using the full template bank, TESLA template bank, and TESLA-X template bank,




Figure 11.4.5: (Top panel) The sensitive range v.s. FAR for MS220508a-alike signals using the full template bank (black), TESLA template bank (red), and TESLA-X template bank (blue), respectively. The shaded band for each curve represents the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. We note that both the TESLA and TESLA-X bank improves in terms of sensitivity towards MS220508a-alike (lensed) signals, but the curve representing the TESLA-X bank (blue) is above that of the TESLA bank (red), meaning that the TESLA-X method outperforms the TESLA method. In particular, notice that the improvement is most significant in the region with $\text{FAR} \geq 10^{-7}$ Hz, which represents the “sub-threshold” region. This further proves that the TESLA-X method is better than the traditional TESLA method.

respectively (see Figure 11.4.5), we can see that the TESLA-X method gives a bigger improvement in terms of sensitivity towards MS220508a-alike (lensed) signals than the traditional TESLA method. In particular, we note that the improvement is most significant in the region with $\text{FAR} \geq 10^{-7}$ Hz, which represents the “sub-threshold” region. This means that TESLA-X outperforms TESLA in finding possible sub-threshold strongly lensed gravitational waves.

11.5 Concluding Remarks

The detectability of strongly lensed gravitational wave (GW) signals can be compromised as they may experience demagnification, falling below the detection threshold. In a previous study, we introduced the TESLA method [223] to extract strongly

lensed sub-threshold GW signals by constructing a target template bank based on the information from detected events. This Chapter presents an enhanced approach called TESLA-X, which aims to improve the search sensitivity for strongly lensed sub-threshold GW signals.

The TESLA-X methodology begins with generating simulated lensed injections using posterior samples obtained from parameter estimation. These injections are then subjected to an injection campaign using GstLAL, with a threshold set at a FAR of 3.385×10^{-7} Hz, ensuring statistical significance. Subsequently, the TESLA-X method utilizes the retrieved injections to create a targeted population model and a densely distributed template bank. Finally, the data is reprocessed to identify potential sub-threshold lensed candidates using the new population model and template bank.

The targeted population model is constructed through several steps. Initially, a Gaussian kernel density estimation (KDE) function is applied to the retrieved injections' source parameters $\vec{\gamma}$. This KDE function is then transformed into a probability density function, representing the parameter space of our potential strong-lensed sub-threshold GW signals [166]. Additionally, we generate the targeted template bank by defining a boundary within the retrieved injections. This process involves discarding templates from the original full template bank outside of a predefined boundary that encloses the target parameter space while keeping all templates within the boundary. This results in a densely distributed template bank targeting a smaller region. Compared to the TESLA method, this refined template bank ensures the recovery of any lensed counterparts with less than a 3% SNR loss. Notably, the FARs and likelihood values (\mathcal{L}) do not provide information about the likelihood of being lensed counterparts to the target event. Instead, they serve as metrics to differentiate between noise triggers and astrophysical signals.

In this study, we conducted a simulation campaign to evaluate and compare the search sensitivity of three methods: the full template bank search, the TESLA method, and the TESLA-X method. Our analysis utilized a stream of mock data, where Gaussian noise was introduced along with a significant GW event labeled MS220508a. Our findings demonstrate that the TESLA-X method outperforms the other approaches by successfully recovering more lensed injections specifically targeting MS220508a. Furthermore, we observed a maximum of $\sim 10\%$ increase in the search sensitivity through an examination of the sensitive range versus FAR plot. An increase in sensitive range means that we can see signals coming from further

sources, and equivalently, we can see weaker signals according to equation 11.1. These results substantiate the superior performance of the TESLA-X method in enhancing the search capabilities for gravitational lensing phenomena. TESLA-X still carries the same problem as TESLA in terms of high computational cost¹⁰ for running the injection campaign for each target super-threshold gravitational-wave event. As future work, we will explore ways to improve TESLA-X's computational efficiency. We will also look into possible parameters other than chirp mass and effective spin to represent our interested population (See Appendix).

¹⁰Running TESLA/TESLA-X for a single super-threshold event can take $O(\text{weeks})$.

APPENDIX

11.A Choosing a suitable pair of parameters to construct the Gaussian KDE

In the main text, we construct the Gaussian KDE in the chirp mass \mathcal{M}_c - effective spin χ_{eff} parameter space. It has also been suggested that the component masses m_1 - m_2 parameter space may be a better choice. In this appendix, we present some preliminary investigation results for this question.

Re-doing the TESLA-X analysis for MS220508a with the Gaussian KDE constructed in the m_1 - m_2 space

In this subsection, we redo the TESLA-X analysis with the Gaussian KDE constructed in the component masses m_1 - m_2 parameter space instead of the chirp mass \mathcal{M}_c - effective spin χ_{eff} parameter space.

Constructing the TESLA-X bank and population model

We start with the injection run results we obtained previously for the same event. First, Figure 11.A.1 plots the rung-up templates on the component masses m_1 - m_2 parameter space, with the colors of the markers being the network SNRs of the associated injections.

Note that although we are working in a different parameter space, templates in the TESLA template bank will remain unchanged because we are simply keeping all the templates that are rung up by the injections during the injection campaign to form the TESLA bank. For completeness, we also plot the **same** TESLA bank in the main text for MS220508a on the component masses m_1 - m_2 parameter space in Figure 11.A.2 As before, templates in the full template bank are plotted in grey, and those in the TESLA template bank are plotted in orange.

For TESLA-X, we use the rung-up templates in the injection campaign to construct a Gaussian Kernel Density Estimation (KDE) function $\text{KDE}(\vec{\gamma})$ with the associated injections' network SNRs as the weights. As before, γ represents the template parameters used to label the templates. Here, $\vec{\gamma} = \{m_1, m_2\}$. Then, we evaluate the probability density at each template point in the full template bank, and plot the results on a contour map. Figure 11.A.3 shows the contour map for the Gaussian KDE function estimated for MS220508a in the m_1 - m_2 space.

We again take the lowest contour level (0.00002) as a boundary, and use all the




Figure 11.A.1: Rung-up templates by the simulated lensed injections of MS220508a in the injection campaign, plotted on the component masses m_1 - m_2 parameter space. The colors of the markers represent the network SNRs of the associated injections.

templates that fall within this boundary to construct a TESLA-X targeted template bank in the m_1 - m_2 space. For simplicity, we will call the TESLA-X bank created based on the KDE constructed in the m_1 - m_2 space as the “TESLA-X (component mass) bank”, and the TESLA-X bank in the main text that is created based on the KDE constructed in the $\mathcal{M}_c - \chi_{\text{eff}}$ space as the “TESLA-X (main) bank”. Figure 11.A.4 shows the templates in the full template bank (in grey) and in the “TESLA-X (component mass) bank” (in red) in the m_1 - m_2 space. The yellow curve marks the lowest contour level we obtained for the estimated Gaussian KDE function we showed in Figure 11.A.3. There are 159043 templates in the “TESLA-X (component mass) bank”.

Similar to the main text, we use the estimated KDE function for MS220508a and the “TESLA-X (component mass) bank” to construct a targeted population model, following the same steps detailed in 11.3 and [166].




Figure 11.A.2: The templates in the full template bank (in grey) and in the TESLA reduced template bank (in orange), plotted in the component masses m_1 - m_2 space.

Performing a re-filtering to recover the lensed injections with the “TESLA-X (main) bank” and “TESLA-X (component mass) bank”

To compare the effectiveness of the “TESLA-X (main) bank” and “TESLA-X (component mass) bank” for the case of MS220508a, we perform an additional search using the “TESLA-X (component mass) bank” together with the targeted population model over the same stretch of mock data described in the main text to try to recover the same set of lensed injections used in the injection campaign. Table 11.A.1 summarizes the findings. We can see that “TESLA-X (component mass)” is performing even worse than the full template bank, with a reduction in recovered lensed injections by -0.3% . This should not be surprising, however, given that the “TESLA-X (component mass) bank” has more templates than the “TESLA-X (main) bank” (almost six times more templates). This results in a much higher trials factor and, hence, a larger noise background, making it harder to recover the simulated lensed injections.

Furthermore, if we look at the sensitive range v.s. FAR for MS220508a-alike signals




Figure 11.A.3: The contour map of the Gaussian Kernel Density Estimation function obtained for MS220508a in the $m_1 - m_2$ space. The colors represent the probability density estimated by the KDE function.

Table 11.A.1: Number of injections found during the search of mock data using the general template bank, “TESLA-X (main) bank” and “TESLA-X (component mass) bank” respectively.

Injections	General	TESLA-X (main)	TESLA-X (component mass)
Total	7815	7815	7815
Found	6151	6297	6132
Found % change	-	+2.37%	-0.3%

using the full template bank, “TESLA-X (main) bank” and “TESLA-X (component mass) bank” respectively (see Figure 11.A.5), we can see that the “TESLA-X (component mass) bank” is performing even worse than the full template bank. Therefore, in the case of MS220508a, chirp mass \mathcal{M}_c - effective spin χ_{eff} would be the better pair of parameters to construct the Gaussian KDE in the TESLA-X method.




Figure 11.A.4: The templates in the full template bank (in grey) and in the “TESLA-X (component mass) bank” (in red), plotted in the m_1 - m_2 space. The “TESLA-X (component mass) bank” is constructed by keeping templates that fall within the lowest contour level from the Gaussian KDE function shown in Figure 11.A.3, i.e. 0.00002, plotted as a yellow curve in this figure. For easy comparison, we also plot the templates in the TESLA reduced template bank in the same figure (in blue).

Analyzing another mock event MS220510ae with the TESLA-X method

MS220508a is a relatively low chirp mass event ($\sim 45.0M_\odot$). It will be useful to redo the above investigation about the choice of parameters to construct the Gaussian KDE in the TESLA-X method with a different mock event that lives in a different region in the parameter space.

In this subsection, we repeat the analysis in the previous subsection with another mock superthreshold gravitational-wave event MS220510ae. Settings for the mock data are identical to those in the main text. Details about the source parameters of the injected gravitational wave are listed in Table 11.A.2.

We follow the same steps in the main text to perform Bayesian parameter estimation on MS220510ae, and use the posterior samples obtained to generate a set of lensed injections. We then conduct an injection campaign using the full template bank to




Figure 11.A.5: (Top panel) The sensitive range v.s. FAR for MS220508a-alike signals using the full template bank (black), “TESLA-X (main) bank” (red), and “TESLA-X (component mass) bank” (blue) respectively. The shaded band for each curve represents the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. We note that while the “TESLA-X (main) bank” improves in terms of sensitivity towards MS220508a-alike (lensed) signals, “TESLA-X (component mass) bank” is performing worse than the full template bank, mainly due to the significant increase in trials factors caused by the large number of templates.

search for the lensed injections. The injection run results are then used to construct the “TESLA-X (main) bank” and “TESLA-X (component mass) bank”, and their corresponding targeted population models. Out of the 8099 lensed injections, the full template bank manages to recover 5047. For simplicity, we only show the Gaussian KDE constructed for the “TESLA-X (main) bank” (in the \mathcal{M}_c - χ_{eff} space) and for the “TESLA-X (component mass) bank” (in the m_1 - m_2 space) in Figure 11.A.6 and 11.A.7 respectively.

As before, we construct the “TESLA-X (main) bank” and “TESLA-X (component mass) bank” by keeping all templates that fall within the boundary defined by the least significant contour in the respective Gaussian KDE constructed. For the case of MS220510ae, the “TESLA-X (main) bank” contains 15655 templates, and the “TESLA-X (component mass) bank” contains 21676 templates. Notice that the “TESLA-X (component mass) bank” for MS220510ae, unlike in the case for

Table 11.A.2: Information of the injected super-threshold gravitational-wave signal MS220510ae in the simulation campaign. All properties reported here are measured in the detector frame.

Properties	Injected super-threshold signal
UTC time	May 10 2022 18 : 38 : 17
GPS time	1336243115.828
Distance (Mpc)	10066.811
Primary mass m_1^{det}	$166.90M_{\odot}$
Secondary mass m_2^{det}	$117.04M_{\odot}$
Dimensionless spins	$\chi_{1x} = 0.635, \chi_{1y} = 0.233, \chi_{1z} = 0.660,$ $\chi_{2x} = 0.006, \chi_{2y} = 0.087, \chi_{2z} = 0.029,$
Right ascension α	5.69
Declination δ	1.45
Inclination ι	5.99
Polarization Ψ	2.92
Waveform	IMRPhenomXPHMpseudoFourPN

MS220508a, is larger than the “TESLA-X (main) bank” by merely ~ 5000 templates.

We then use the two TESLA-X banks and their respective population models to try recovering the same set of lensed injections used in the injection campaign. Table 11.A.3 summarizes the findings. Completely different from the case of MS220508a in the main text, we can see that “TESLA-X (component mass)” is performing even better than the full template bank, while the “TESLA-X (main) bank” is performing worse than the full template bank. This is mainly because spins are poorly measured in general for high-mass systems, and hence the “TESLA-X (main) bank”, with inaccurate constraint in the spin dimension, are more likely to lose signals than the “TESLA-X (component mass) bank”, which does not impose any constraints on spins. For completeness, we also show the sensitive range v.s. FAR for MS220510ae-alike signals using the full template bank, “TESLA-X (main) bank” and “TESLA-X (component mass) bank” respectively in Figure 11.A.8.

Conclusion

From the two investigations, we note that “chirp mass - effective spins” is a better parameter pair to be used in constructing the Gaussian KDE for the TESLA-X method, if the target event is a low chirp mass system. On the other hand, component masses “mass 1 - mass 2” will be a better parameter pair to be used in constructing the Gaussian KDE for the TESLA-X method, if the target event is a high chirp mass




Figure 11.A.6: The contour map of the Gaussian Kernel Density Estimation function obtained for MS220510ae in the \mathcal{M}_c - χ_{eff} space. The colors represent the probability density estimated by the KDE function.

Table 11.A.3: Number of injections found during the search of mock data using the general template bank, “TESLA-X (main) bank” and “TESLA-X (component mass) bank” respectively for the mock event MS220510ae.

Injections	General	TESLA-X (main)	TESLA-X (component mass)
Total	8099	8099	8099
Found	5047	4715	5886
Found % change	-	-6.58%	16.6%

system. We note that these findings are entirely expected: While chirp masses and spins are well-measured in low-mass systems, they are not well-measured in high-mass systems since their signals have much fewer in-band cycles. As future work, we will investigate how we can improve the construction of the Gaussian KDE functions for the TESLA-X method. For instance, we can see that the Gaussian KDEs can exceed the physical boundary of the original template bank (See the yellow boundary in 11.A.4). Regions near the physical boundary of the original




Figure 11.A.7: The contour map of the Gaussian Kernel Density Estimation function obtained for MS220510ae in the $m_1 - m_2$ space. The colors represent the probability density estimated by the KDE function.

template bank may require special care. We may also explore the idea of first constructing and applying a separate KDE to “flatten” the TESLA-X bank (i.e., to construct a uniform mass model) before applying the Gaussian KDE representing the targeted population model.




Figure 11.A.8: (Top panel) The sensitive range v.s. FAR for MS220510ae-alike signals using the full template bank (black), “TESLA-X (main) bank” (red), and “TESLA-X (component mass) bank” (blue) respectively. The shaded band for each curve represents the corresponding 1-sigma region. (Bottom panel) The corresponding percentage changes in sensitive range v.s. FAR for the different banks. We note that, unlike the case of MS220508a in the main text, the “TESLA-X (component mass) bank” improves in terms of sensitivity towards MS220510ae-alike (lensed) signals, “TESLA-X (main) bank” performs worse than the full template bank. This is mainly because spins are typically poorly measured for high chirp mass systems.

RE-RANKING POSSIBLE SUB-THRESHOLD LENSED CANDIDATES WITH LENSING LIKELIHOODS

Note: This chapter will be turned into a short author publication:

Alvin K. Y. Li et al., “TESLA-X+: Targeted search method for sub-threshold strongly lensed gravitational waves with lens-model-based ranking statistics”, In Preparation.

12.1 Introduction

The GstLAL-based TESLA / TESLA-X pipeline [222, 223] has been one of the key pipelines for the LVK’s collaboration to search method for possible sub-threshold lensed counterparts to known superthreshold gravitational waves [23, 32]. By doing injection campaigns, TESLA / TESLA-X narrows down the search parameter space and reduces the size of the template bank to lower the noise background. With the help of a targeted population model, they help to improve the ranking statistics of possible sub-threshold lensed gravitational-wave candidates. While the number of templates in targeted template banks generated through the TESLA-X pipeline is typically much lower than that used in searches for gravitational waves in general ¹, targeted searches using the reduced template banks can still return $O(10)$ of possible lensed candidates per event. To do follow-up on these candidates to see how likely they are lensed counterparts to the target super-threshold gravitational-wave event, one will first need to obtain the posterior probability distribution that gives the best estimates of the source parameters of the target using the Bayesian parameter estimation analyses described in [48, 49, 305, 356]. For sub-threshold candidates that typically have lower signal-to-noise ratios (SNRs), conducting parameter estimation analyses are extremely computationally expensive and time-consuming. It is, therefore, vital to develop ways to discard candidates returned from each TESLA-X analysis effectively that are unlikely lensed counterparts to the target event, before passing them for follow-up analysis.

This chapter introduces a new ranking statistic for the TESLA-X pipeline incorporat-

¹ For reference, the template bank used by the LVK collaboration in O3 to search for gravitational waves has $O(10^6)$ templates, whereas a typical TESLA-X targeted bank has $O(10^3 - 10^4)$ templates.

ing lens model information. This new ranking statistic, named “lensing likelihood”, evaluates how likely a candidate returned from the TESLA-X pipeline is a lensed counterpart to the target superthreshold event, assuming that the candidate is an actual gravitational-wave event. It is important to note that the lensing likelihood is lens-model-based (i.e. its value can vary for different lens models). The possible use cases of the lensing likelihood will be discussed later.

This chapter is organized as follows: Section 12.2 describes the motivation for implementing the lensing likelihood. In Section 12.3, we describe how the lensing likelihood is constructed and evaluated. Section 12.4 presents a working example of how the lensing likelihood can help to prioritize candidate-follow-up for the TESLA-X pipeline. Finally, in Section 12.5, we summarize our findings and discuss potential future work to further enhance the usefulness of the newly implemented lensing likelihood.

12.2 Motivation and Background

Strong gravitational lensing can produce copies of gravitational waves from the same source. These copies have the same intrinsic parameters (hence the same waveform morphology) but with different amplitudes (due to different amplification factors applied) and phases (due to different Morse phase factors applied). Some of these copies are de-amplified compared to the not-lensed gravitational wave because the amplification factor is < 1 , and may become “sub-threshold”². GstLAL-based TESLA-X [222, 223] (See also Chapter 11) has been the flagship LVK matched-filtering-based search pipeline to search for strongly-lensed sub-threshold lensed gravitational-wave counterparts to confirmed superthreshold gravitational waves from LVK’s first three observing runs [23, 32, 186]. In this section, we will give a brief review of the TESLA-X pipeline, and motivate the implementation of the lensing likelihood, which will be discussed in the next section.

The TESLA-X search pipeline: A brief recap

TESLA-X aims to search for possible sub-threshold lensed counterparts to known superthreshold gravitational waves. These possible new candidates are not identified as gravitational waves in the general search by the LVK collaboration because a large template bank spanning a vast region in the search parameter space is used. As more templates are used in a search, more candidates (triggers), whether they are

² As explained in previous chapters, “sub-threshold” signals are low-amplitude signals that are indistinguishable from the noise background.

real signals or noise, will be generated. Because of the high trials factor, the noise background generated in a general search for gravitational waves will be large, hence making it more difficult for low-amplitude signals, e.g. possible sub-threshold lensed gravitational waves, to be distinguished from noise and be identified as signals. TESLA-X strategically reduces the nuisance noise background while keeping the targeted foreground constant to uncover possible gravitational waves with weaker amplitudes.

For a given superthreshold gravitational-wave event, TESLA-X first considers the posterior probability distribution that provides the best estimates of the source parameters of the target using BILBY, a Bayesian parameter estimation pipeline [48, 49, 305, 356]. Under the strong lensing hypothesis, we expect possible sub-threshold lensed counterparts to the target superthreshold gravitational-wave event to have similar intrinsic parameters. Therefore, it is natural to consider localizing the search parameter space to regions consistent with the target's posterior parameter space. However, we note that the posterior distributions obtained for the target event are only valid for one noise realization. Due to noise fluctuations in the data, weaker signals (near or sub-threshold) can be registered as triggers by templates with parameters very different from those of the posterior samples of the target event [222, 223].

To account for both consistency in intrinsic parameters with the target event, and noise non-stationarity, TESLA-X takes the posterior samples of a target superthreshold gravitational-wave event from parameter estimation as input, and sorts them in decreasing order of log-likelihood. We aim to determine the region in the parameter space where weaker lensed counterparts can be recovered. These possible lower-amplitude signals will be registered as triggers in searches with lower signal-to-noise ratios ρ (SNR). For each sample, TESLA-X generates 1 injection with identical parameters (including luminosity distance to the source D_L and sky location, i.e. right ascension α and declination δ) as the sample, and 9 extra injections with increasing D_L to mimic the effect of de-magnification due to strong lensing on the amplitudes of the possible sub-threshold signals, according to the relation $D_L \propto \rho^{-1}$.

The set of lensed injections is injected into real data. Then, an injection campaign is performed using GstLAL with a general, full template bank to try to recover the injections (As before, an injection is recovered if the false-alarm-rate FAR of the associated trigger is ≤ 1 in 30 days.). Recovered injections with their associated templates are then used to generate a Gaussian Kernel Density Estimation (KDE)

$p(\vec{\theta})$ on the search parameter space represented by $\vec{\theta}$ ³ with the recovered SNR of the injections as weights. Templates that live within the lowest contour of the KDE are kept to construct a targeted template bank. The KDE is further used to generate a targeted population model [166, 222] that describes the expected distribution of possible (sub-threshold) lensed counterparts to the target super-threshold gravitational wave.

Finally, TESLA-X uses the GstLAL search pipeline again to search through all possible data with the generated targeted template bank and population model. GstLAL will output a list of possible candidate events ranked according to their assigned ranking statistics, including the likelihood ratios \mathcal{L} and false-alarm-rates FARs. These ranking statistics assigned by GstLAL, however, only represents how likely the candidate is an actual gravitational-wave event as compared to noise, i.e. it does not include any additional information regarding how likely the candidate event is a lensed counterpart to the target event, compared to being an independent event.

Candidates that satisfy a pre-determined FAR threshold are then passed on for follow-up analyses, which require parameter estimation to be first done on each new candidate. Since we expect to see many sub-threshold candidates with lower SNRs, it is, in general, more computationally costly and time-consuming to conduct parameter estimation for them.

In the LVK’s search for lensing signatures in full O3 data [32], the TESLA pipeline reported 2372 new sub-threshold candidates that passes the usual 1 in 30 days FAR threshold from targeted searches for 39 superthreshold gravitational-wave events [26]. This corresponds to approximately 60 new candidates per event. In the paper, we set a stricter threshold on the FAR and only follow up on the most promising $O(1)$ candidates due to the expensive computational costs. However, we note that the FAR assigned to each candidate only tells us how likely they are gravitational-wave events as compared to noise. Lensing information (i.e., how likely these candidates are lensed counterparts to the target event) has not been considered. To better trim the list of candidates and only follow up on those that are more likely to be an actual gravitational-wave event, and more likely to be lensed counterparts to the target event, we need a second alternative ranking statistic that accounts for

³ $\vec{\theta}$ labels the intrinsic parameters of the templates in the search parameter space. They can be different sets of parameters, e.g. “chirp mass \mathcal{M}_c - effective spin χ_{eff} ” or “component mass 1 m_1 - component mass 2 m_2 ”, depending on the region of parameter space we are probing.

the consistency of the candidates' parameters with the lensing hypothesis. In the next section, We will give a detailed discussion about this new statistic, the lensing likelihood.

12.3 The Lensing Likelihood

In gravitational lensing, the primary observables are the relative time delays and magnifications between the pair of lensed events [259]. Given a lens model, one can simulate lensed gravitational-wave events and model the expected distributions of the relative time delays $p(\Delta t|\text{lensed})$ and magnifications (Magnification ratios) $p(\mu_{\text{rel}}|\text{lensed})$ of the possible lensed event pairs. We will use the joint probability distribution $p(\Delta t, \mu_{\text{rel}}|\text{lensed})$ instead of the factorized probability distribution. We can repeat the exercise for independent (not lensed) gravitational-wave events, and obtain the joint probability distribution of their relative time delays and magnifications $p(\Delta t, \mu_{\text{rel}}|\text{not lensed})$ [259]. Figure 12.3.1 shows an example of the Kernel Density Estimations of the distribution of log relative time delay and magnification ratios for the simulated lensed gravitational waves (green contours and pale blue sample points) for the SIE-double model and not-lensed, independent gravitational waves (black contours with grey sample points). We can see that the KDE for not-lensed, independent gravitational waves spans a larger range of relative magnifications and relative time delays solely because these gravitational-wave events are randomly distributed, in contrast to lensed gravitational waves under a given lens model. Note that here we are only considering galaxy-scale lenses that typically have lens masses ranging from $10^6 M_\odot$ to $10^{13} M_\odot$ (See, for example, Table 9.5.1 in Chapter 9). For galaxy-cluster-scale lenses with lens masses $\geq 10^{13} M_\odot$, lensed gravitational waves can have larger relative time delays (i.e. of the order of years), with larger absolute magnifications.

Since lensed and not lensed gravitational waves have different distributions for relative time delays and relative magnifications. one can construct the proposed “lensing likelihood” with these joint distributions as

$$\mathcal{L}^{\text{lensing}} = \frac{p(\Delta t, \mu_{\text{rel}}|\text{lensed, signal})}{p(\Delta t, \mu_{\text{rel}}|\text{not lensed, signal})}. \quad (12.1)$$

For candidates returned from the search pipeline, relative magnification can be defined as the ratio of the candidate's SNR to that of the target event. Note that we have implicitly assumed that the candidates under consideration are real gravitational-wave signals.




Figure 12.3.1: The Kernel Density Estimations of the distribution of log relative time delay and magnification ratios for the simulated lensed gravitational waves (green contours and pale blue sample points) for the SIE-double model and not-lensed, independent gravitational waves (black contours with grey sample points). We can see that the KDE for not-lensed, independent gravitational waves spans a larger range of relative magnifications and relative time delays solely because these gravitational-wave events are randomly distributed, in contrast to lensed gravitational waves under a given lens model.

In contrast to the likelihood assigned by search pipelines that evaluate how likely a candidate is a gravitational-wave signal as compared to noise, the lensing likelihood evaluates how likely a candidate is a lensed counterpart to a given target event, as compared to being completely independent events, by considering the relative time delay and magnification between the two events under consideration. The higher the lensing likelihood, the more consistent the lensing observables (i.e. relative time delay and magnification) of the candidate are with the lens model under consideration, and thus, the more likely the candidate is a lensed counterpart to the target event under consideration.

There are two caveats to the interpretation of the lensing likelihood: (1) The lensing likelihood assumes the candidate is an actual gravitational-wave event. This means

that the lensing likelihood only assesses the consistency with the strong lensing hypothesis of the candidates' relative time delay Δt and relative magnification μ_{rel} compared to the target event, regardless of whether the candidate is an actual signal or noise. (2) The lensing likelihood heavily depends on the lens model under consideration. Different lens models have different expected distribution for the relative time delays and magnifications between lensed signals, and hence, the lensing likelihood for the same candidate can change when one varies from one lens model to another.

12.4 Working example

This section uses the mock lensed event MGW220111a introduced in [223] as a working example for the lensing likelihood implementation in TESLA. As a quick recap, a pair of lensed events (GPS time: 1325932493 and 1326029051), one superthreshold and one subthreshold, is injected into a 28-hour-long data stream with Gaussian noise recolored with O3a characteristic power spectral densities (PSDs). The two signals are image type 1 and type 2 signals, respectively. Detailed setup and parameters of the injected events can be found in Section IV of [223], and Section 7.3 of Chapter 7.

In the TESLA targeted search for the mock superthreshold event MGW220111a in [223], we can uprank the missing sub-threshold lensed counterpart (GPS time: 1326029051) and retrieve it as a rank-3 candidate, followed by the target event MGW220111a (rank-1 candidate) and a noise trigger (rank-2 candidate). Table 12.4.1 shows the detailed information of the top three candidates from the TESLA targeted search. Should this situation be encountered in a real analysis, we will

Search results	Superthreshold signal	Noise trigger	Subthreshold signal
GPS time	1325932493	1326011224	1326029051
$\log_{10} \Delta t$ (days)	-	-0.04	0.04
Rank	1	2	3
FAR (Hz)	5.37×10^{-21}	1.18×10^{-5}	4.27×10^{-5}
$\ln \mathcal{L}$	48.63	13.53	12.13
Network SNR ρ_{network}	12.20	7.12	7.60
$\log_{10} \mu_{\text{rel}}$	-	0.23	0.21

Table 12.4.1: Top three candidates from the TESLA targeted search for the mock lensed event MGW220111a.

have to follow up on both the rank 2 and rank 3 triggers, because we do not know a

priori whether they (1) correspond to a real gravitational wave, and (2) are lensed counterparts to the target event (i.e. rank 1 trigger).

Following [259], we simulated lensed gravitational waves based on the singular isothermal ellipsoid (SIE) lens model. The SIE lens model allows for systems of 2 lensed images (also known as “doubles”) and systems of 4 images (also known as “quads”). For our example, we will only consider the SIE-double model. Based on the simulated lensed gravitational waves, we generate Gaussian Kernel Density Estimation (KDE) for the relative time delays $\log_{10} \Delta t$ (in days) and magnification ratios ($\log_{10} \mu_{\text{rel}}$) between the lensed images. We repeat the exercise with simulated, not-lensed, independent gravitational waves. Then, for both the rank 2 and rank 3 triggers listed in Table 12.4.1, we evaluate the log lensing likelihood $\ln \mathcal{L}^{\text{lensing}}$. Table 12.4.2 shows the evaluated results. In Figure 12.4.1 we also plot the relative time delays and magnification ratios for the rank 2 (in red) and rank 3 (in blue) triggers, on top of the KDE distributions for lensed gravitational waves (green contours) and not-lensed, independent gravitational waves (filled contours). We can see that the lens-

Log Lensing	Rank 2	Rank 3
Likelihood	trigger	trigger
$\ln \mathcal{L}^{\text{lensing}}$	3.80	3.98

Table 12.4.2: Log Lensing Likelihoods for the rank 2 and 3 triggers evaluated based on the SIE-double lens model.

ing likelihood of the rank 3 trigger, corresponding to the real sub-threshold lensed gravitational-wave counterpart to the mock superthreshold event MGW220111a, has a “marginally-larger” log lensing likelihood ($\ln \mathcal{L}^{\text{lensing}} = 3.98 > 3.80$) than that of the rank 2 trigger, corresponding to a noise trigger. In Figure 12.4.1, we can see that the rank 3 trigger’s (in blue) lensing observables (relative time delay and magnification) is slightly closer to the lensed gravitational wave distribution (green contours), compared to those of the rank 2 trigger (in red). While being a marginal case, this still demonstrates that the lensing likelihood has the potential to uprank candidates that are more likely to be a lensed counterpart to the target event under consideration. However, it should be noted that the lensing likelihood calculation is based on the assumption that the candidate under consideration is a real gravitational wave. This is not the case in this example, where the rank 2 trigger is a noise trigger. In addition, the lensing likelihood is highly dependent on the lens model under consideration, and may vary vastly should a different lens model be




Figure 12.4.1: The log relative time delay and magnification ratios of the rank 2 trigger (in red) and rank 3 trigger (in blue) respectively. The Kernel Density Estimations of the distribution of log relative time delay and magnification ratios for the simulated lensed gravitational waves (green contours) and not-lensed, independent gravitational waves (filled contours) are also plotted. We note that the rank 3 trigger, corresponding to the real sub-threshold lensed counterpart to MGW220111a, is slightly closer to the lensed KDE distribution than the rank 2 noise trigger, resulting in a slightly higher log lensing likelihood.

considered. Therefore, we remind readers that the lensing likelihood, at the current stage, should be treated with caution and only be used as reference information.

12.5 Conclusion and future work

TESLA-X has been the LVK collaboration’s key pipeline to search for possible sub-threshold lensed counterparts to known superthreshold gravitational-wave events. As of the O3 analysis, the TESLA-X pipeline can return 60, possibly lensed sub-threshold candidates per superthreshold event. Due to the expensive computational costs, we trim down the candidate list by imposing a high threshold on the false-alarm-rate FAR assigned to the candidates by the search pipeline. However, as mentioned in [222, 223], the ranking statistics assigned by the search pipeline to

each candidate only consider how likely the candidate is an actual gravitational-wave signal compared to noise. It contains no information about how likely the candidate is a lensed counterpart to the target event under consideration. Therefore, a pure cut in FAR to reduce the number of candidates to follow up with may potentially lead to real lensed gravitational waves being discarded.

In this chapter, we introduce an alternative ranking statistic for the TESLA-X pipeline that incorporates lens model information. The new ranking statistic, or “lensing likelihood”, evaluates how likely a candidate returned from the TESLA-X pipeline is a lensed counterpart to the target superthreshold event, assuming that the candidate is an actual gravitational-wave event. The lensing likelihood takes in the relative time delay and magnification ratios between the sub-threshold lensed candidate and the target event, and compare the probability of obtaining this set of observables under the hypothesis that they are lensed counterparts to each other, as compared to under the hypothesis that they are not-lensed, independent events.

We demonstrate how the lensing likelihood can potentially uprank real lensed sub-threshold gravitational waves with the mock gravitational-wave event MGW220111a case as an example. However, we note the following caveats to the interpretation of the lensing likelihood: (1) The lensing likelihood is evaluated based on the assumption that the candidate is an actual gravitational-wave event. (2) The lensing likelihood depends heavily on the lens model under consideration.

There are several possible use cases for the lensing likelihood proposed in this chapter. (1) The lensing likelihood can be used as a separate ranking statistic to uprank / downrank candidates, (2) We can also multiply the lensing likelihood to the search pipeline’s likelihood that addresses how likely the candidate is a signal as compared to noise, then use the combined ranking statistic to uprank / downrank candidates. We may as well just use the new ranking statistic as a rank / label, but we can also run simulation campaigns to understand the distribution of lensing likelihoods for lensed gravitational-wave events and independent gravitational-wave events to give a better statistical interpretation of the new ranking statistic values. In the future, we can also consider other lens models and different types of lenses (e.g. galaxies and galaxy clusters). Since the calculations are done after the searches as a post-processing step, the additional computational cost of considering more lens models and lens types will be minimal.

Part III

Summary

SUMMARY AND FUTURE WORK

13.1 Summary of Thesis

The detection of gravitational waves by the LIGO-Virgo-KAGRA (LVK) collaboration has opened a new window for us to study the universe. We have detected gravitational waves from compact binary coalescences, which are mergers of binary compact objects: black holes and/or neutron stars. To continue pushing the boundaries to explore new science via gravitational waves, we need to enhance both the hardware, e.g. improving the gravitational-wave detectors' sensitivities and building more ground-based detectors to expand the detector network; as well as software, e.g. developing better search pipelines for gravitational waves and data analysis techniques to probe gravitational-wave sources with the broadest range of parameters and from the furthest distances. With these efforts, we aim to populate our catalog of gravitational-wave events maximally not restricted to only known and confirmed sources like compact binary coalescences, but also, other types of gravitational waves, including gravitational-wave bursts from core-collapse supernovae, continuous gravitational waves from spinning non-axisymmetric neutron stars, and stochastic gravitational-wave background from sources that are not individually resolvable. These detections allow us to explore many essential and exciting topics in astrophysics that can foster our understanding of our universe's fundamental laws and evolution. For instance, gravitational waves can be deflected by curvature in spacetime caused by massive intervening objects like galaxies and galaxy clusters as electromagnetic (EM) waves are, an effect known as gravitational lensing. Successful detection of lensed gravitational waves will be a tremendous discovery for the (astro)physics community. It will also be a pivotal pathway to study the higher-redshift early universe and, ultimately, cosmic evolution.

In this thesis, I focus on improving the search sensitivity of the matched-filtering based search pipeline GstLAL, and searching for gravitationally-lensed gravitational waves. Here, I will briefly summarize the content of my thesis. Beginning with Chapter 2, I gave a very brief overview of the key ideas and concepts in General Relativity, followed by the standard derivations of linearized gravity in the weak-field limit, and how gravitational waves emerge from the resulting Einstein's field

equations. I also briefly introduced the Advanced LIGO gravitational-wave detectors and derived the antenna pattern functions associated with the LVK collaborations' L-shaped detectors. Then, in Chapter 3, I gave a top-level review of the current framework of the LVK collaboration to search for gravitational waves from compact binary coalescences, including data calibration, searches for gravitational waves in calibrated data, and Bayesian parameter estimation analysis for gravitational-wave events. In Chapter 4, I explained in detail how matched-filtering pipelines search for gravitational waves in general. I then focus on GstLAL, one of the flagship pipelines of the LVK collaboration to search for gravitational waves from compact binary coalescences. Alongside the overall structure of the GstLAL pipeline, I have also selected and explained several key terms in the likelihood ratio calculation for triggers in the GstLAL pipeline. With proper modifications, these chosen terms will allow us to enhance the search sensitivity of GstLAL further. For instance, in Chapter 5, I introduced a method to effectively incorporate statistical data quality information from the machine-learning-based pipeline iDQ [149, 176] into GstLAL's likelihood ratio calculations. This method aims to (partially) resolve the problem of current GstLAL noise background estimation, which does not track time dependence nor data quality. Improvements to the other highlighted terms are left as future work (See the next section).

In the second part of my thesis, I switched my focus towards gravitational lensing of gravitational waves. In Chapter 6, I gave an overview of gravitational lensing in general and gravitational lensing of gravitational waves. I also introduced current efforts made by the LVK collaboration to search for lensing signatures in gravitational-wave data. In Chapter 7, I introduced the targeted sub-threshold search for strongly lensed gravitational waves (TESLA) method. TESLA aims to search for possible sub-threshold (demagnified by gravitational lenses and hence with weaker amplitudes) strongly-lensed counterparts to confirmed superthreshold gravitational waves. To reduce the noise background effectively while keeping the targeted foreground constant, TESLA conducts injection campaigns to narrow the search parameter space, accounting for both the target super-threshold event signal subspace and noise fluctuations in the data. Through a simulation campaign, we have shown that TESLA can boost our search sensitivity towards possible sub-threshold lensed gravitational waves in the data, if they exist. In Chapter 8, 9 and 10, we present results from the LVK collaboration-wide search for gravitationally-lensed gravitational waves from the first-half and entire third observing run (O3). The TESLA analysis to search for possible sub-threshold lensed gravitational waves was included

as one of the key analyses. While we have found no evidence of gravitational lensing in O3 data, we expect that with the detectors' sensitivities improving, upcoming LVK's observing runs will continue to populate our catalog of gravitational-wave detections that will foster our study and understanding of gravitational lensing of gravitational waves. With the help of multi-messenger astronomy and improved analysis methods, we have high hopes of making the first (and many more) detection of gravitationally-lensed gravitational waves in the near future. In Chapter 11 and 12, I further introduced two significant improvements to the TESLA search method. Chapter 11 describes how we can construct a better reduced template bank and a targeted population model to improve our sensitivity towards possible sub-threshold lensed gravitational waves. Chapter 12 demonstrates how we can fine-tune the list of possible lensed gravitational-wave candidates reported by the TESLA search method by the incorporation of lens model information. These improvements are crucial and critical to improve our chance of detecting the first gravitationally-lensed gravitational waves.

13.2 Future work

A wise friend once told me, "There is always room for improvement, and things are not done all at once. They are completed in stages." His words apply to both research work and science development. While I am ending my work as a PhD candidate here, there is still a lot of work left to do to exhaust the potential of gravitational waves to fully understand the Universe.

The TESLA-X+ pipeline can be further improved:

Firstly, the current TESLA-X method arbitrarily defines the region of interest in the search parameter space as enclosed by the least significant contour of the Gaussian Kernel Density Estimation we constructed using the recovered injections from the injection campaign. While we have shown that TESLA-X can recover more lensed injections than the traditional TESLA method and attain a larger sensitivity range, the current approach does discard certain templates that can recover lensed injections, but are outside of the least significant contour of the Gaussian KDE constructed. In the future, we will investigate ways to improve the construction of the targeted population model further. For instance, we may consider first "flattening" the TESLA-X bank (by constructing a uniform mass model) before applying the targeted population model. We may also need to address the problem that the Gaussian KDE can exceed the physical boundary of the original template bank.

Secondly, as discussed in Chapter 12, additional injection campaigns are needed to quantify the significance of the lensing likelihood values reported by the TESLA-X+ pipeline. Currently, they can only be used as reference information because we lack a robust study to understand the distributions of the lensing likelihoods of “background”, not-lensed gravitational waves, compared to that for truly lensed gravitational waves. We will need to conduct more thorough background studies to unleash the full potential of the lensing likelihood.

The lensing likelihood can also be extended to include sky location information of the triggers’ sources. Because the uncertainty in source localization for gravitational waves is much larger than the deviation in source positions of lensed gravitational waves from the same source, we can assume lensed gravitational waves originating from the same source to have the same sky localization. This can be potentially added as an additional parameter to be considered in the lensing likelihood and enhance its power to distinguish possible lensed gravitational waves from not-lensed gravitational waves.

We can also work on the GstLAL search pipeline for more general sky-location-based targeted searches. We want to modify the GstLAL search pipeline to only search for gravitational waves from a specific region in the sky. The use case for sky-location-based targeted searches are not only constrained to searching for strongly-lensed gravitational waves. One can also apply the same methodology to look for gravitational-wave counterparts to EM observations/detections. For instance, a potential use case would be to conduct targeted Gamma-Ray Burst searches [28, 165] / Fast Radio Burst searches [31] that look for gravitational-wave counterparts to detected Gamma-Ray Bursts / Fast Radio Bursts. To do so, however, we must modify the likelihood ratio calculation in the GstLAL search pipeline to include parameters that encode the sky localization of triggers. Additionally, we must develop a (semi-analytic) signal and noise model to include the sky-location-related parameters in the likelihood ratio. Moreover, we will need a rapid method implemented in GstLAL to efficiently localize the source for the many triggers reported by GstLAL. At the same time, we will also need to implement a better “fast-path cut” function to drastically reduce the number of triggers to be followed up without discarding triggers that may be an actual gravitational-wave signal. Should such a “fast-path cut” function be implemented, we will be able to bypass the current clustering process in the GstLAL pipeline and enable the reranking of triggers in the future, potentially removing the need to refiltering data due to the use

of a subset of templates in the full template bank (e.g. the TESLA-X method), or the use of an alternative population model.

We are entering a golden age to pursue gravitational-wave science with multi-messenger astronomy and observational high-energy astrophysics, which requires the worldwide collaboration of the LVK collaboration and EM astronomers and telescopes. There is enormous potential to learn about the early universe through gravitational lensing and multi-messenger astronomy with gravitational waves. The truth is out there for us to discover. . .

BIBLIOGRAPHY

- [1] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015. doi: 10.1088/0264-9381/32/7/074001.
- [2] J. Abadie et al. Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO’s Sixth Science Run and Virgo’s Science Runs 2 and 3. *Phys. Rev. D*, 85:082002, 2012. doi: 10.1103/PhysRevD.85.082002.
- [3] B. P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Class. Quant. Grav.*, 33 (13):134001, 2016. doi: 10.1088/0264-9381/33/13/134001.
- [4] B. P. Abbott et al. The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. *Astrophys. J. Lett.*, 833(1):L1, 2016. doi: 10.3847/2041-8205/833/1/L1.
- [5] B. P. Abbott et al. GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. *Phys. Rev. D*, 93(12):122003, 2016. doi: 10.1103/PhysRevD.93.122003.
- [6] B. P. Abbott et al. Properties of the Binary Black Hole Merger GW150914. *Phys. Rev. Lett.*, 116(24):241102, 2016. doi: 10.1103/PhysRevLett.116.241102.
- [7] B. P. Abbott et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. *Phys. Rev. D*, 95(6): 062003, 2017. doi: 10.1103/PhysRevD.95.062003.
- [8] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017. doi: 10.1103/PhysRevLett.119.161101.
- [9] B. P. Abbott et al. Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA. *Living Rev. Rel.*, 21(1):3, 2018. doi: 10.1007/s41114-020-00026-9.
- [10] B. P. Abbott et al. Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo. *Astrophys. J. Lett.*, 882(2):L24, 2019. doi: 10.3847/2041-8213/ab3800.
- [11] B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9(3):031040, 2019. doi: 10.1103/PhysRevX.9.031040.

- [12] B. P. Abbott et al. GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$. *Astrophys. J. Lett.*, 892(1):L3, 2020. doi: 10.3847/2041-8213/ab75f5.
- [13] Benjamin P. Abbott et al. Effects of waveform model systematics on the interpretation of GW150914. *Class. Quant. Grav.*, 34(10):104002, 2017. doi: 10.1088/1361-6382/aa6854.
- [14] Benjamin P. Abbott et al. Exploring the Sensitivity of Next Generation Gravitational Wave Detectors. *Class. Quant. Grav.*, 34(4):044001, 2017. doi: 10.1088/1361-6382/aa51f4.
- [15] Benjamin P. Abbott et al. A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals. *Class. Quant. Grav.*, 37(5):055002, 2020. doi: 10.1088/1361-6382/ab685e.
- [16] R. Abbott et al. GW190521: A Binary Black Hole Merger with a Total Mass of $150M_{\odot}$. *Phys. Rev. Lett.*, 125(10):101102, 2020. doi: 10.1103/PhysRevLett.125.101102.
- [17] R. Abbott et al. GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses. *Phys. Rev. D*, 102(4):043015, 2020. doi: 10.1103/PhysRevD.102.043015.
- [18] R. Abbott et al. Properties and Astrophysical Implications of the $150 M_{\odot}$ Binary Black Hole Merger GW190521. *Astrophys. J. Lett.*, 900(1):L13, 2020. doi: 10.3847/2041-8213/aba493.
- [19] R. Abbott et al. GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object. *Astrophys. J. Lett.*, 896(2):L44, 2020. doi: 10.3847/2041-8213/ab960f.
- [20] R. Abbott et al. Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo’s third observing run. *Phys. Rev. D*, 104(2):022004, 2021. doi: 10.1103/PhysRevD.104.022004.
- [21] R. Abbott et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. *Phys. Rev. X*, 11:021053, 2021. doi: 10.1103/PhysRevX.11.021053.
- [22] R. Abbott et al. Population Properties of Compact Objects from the Second LIGO–Virgo Gravitational-Wave Transient Catalog. *Astrophys. J. Lett.*, 913(1):L7, 2021. doi: 10.3847/2041-8213/abe949.
- [23] R. Abbott et al. Search for Lensing Signatures in the Gravitational-Wave Observations from the First Half of LIGO–Virgo’s Third Observing Run. *Astrophys. J.*, 923(1):14, 2021. doi: 10.3847/1538-4357/ac23db.

- [24] R. Abbott et al. Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences. *Astrophys. J. Lett.*, 915(1):L5, 2021. doi: 10.3847/2041-8213/ac082e.
- [25] R. Abbott et al. Tests of General Relativity with GWTC-3. *arXiv e-prints*, art. arXiv:2112.06861, December 2021. doi: 10.48550/arXiv.2112.06861.
- [26] R. Abbott et al. Data release for "search for gravitational-lensing signatures in the full third observing run of the ligo-virgo network", 2021. URL <https://zenodo.org/records/10059032>.
- [27] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run — Parameter estimation data release, November 2021. URL <https://doi.org/10.5281/zenodo.5546663>. LIGO Laboratory and Advanced LIGO are funded by the United States National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. The construction and operation of KAGRA are funded by Ministry of Education, Culture, Sports, Science and Technology (MEXT), and Japan Society for the Promotion of Science (JSPS), National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea, Academia Sinica (AS) and the Ministry of Science and Technology (MoST) in Taiwan.
- [28] R. Abbott et al. Search for Gravitational Waves Associated with Gamma-Ray Bursts Detected by Fermi and Swift during the LIGO–Virgo Run O3b. *Astrophys. J.*, 928(2):186, 2022. doi: 10.3847/1538-4357/ac532b.
- [29] R. Abbott et al. Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3. *Phys. Rev. X*, 13(1):011048, 2023. doi: 10.1103/PhysRevX.13.011048.
- [30] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run. *Phys. Rev. X*, 13(4):041039, 2023. doi: 10.1103/PhysRevX.13.041039.
- [31] R. Abbott et al. Search for Gravitational Waves Associated with Fast Radio Bursts Detected by CHIME/FRB during the LIGO–Virgo Observing Run O3a. *Astrophys. J.*, 955(2):155, 2023. doi: 10.3847/1538-4357/acd770.

- [32] R. Abbott et al. Search for gravitational-lensing signatures in the full third observing run of the LIGO-Virgo network. *arXiv e-prints*, art. arXiv:2304.08393, April 2023. doi: 10.48550/arXiv.2304.08393.
- [33] R. Abbott et al. GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run. *Phys. Rev. D*, 109(2):022001, 2024. doi: 10.1103/PhysRevD.109.022001.
- [34] Rich Abbott et al. Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo. *SoftwareX*, (13):100658, 2021. doi: 10.1016/j.softx.2021.100658.
- [35] Rich Abbott et al. Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo. *SoftwareX*, 13:100658, 2021. doi: 10.1016/j.softx.2021.100658.
- [36] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015. doi: 10.1088/0264-9381/32/2/024001.
- [37] F. Acernese et al. Calibration of Advanced Virgo and Reconstruction of the Gravitational Wave Signal $h(t)$ during the Observing Run O2. *Class. Quant. Grav.*, 35(20):205004, 2018. doi: 10.1088/1361-6382/aadf1a.
- [38] F. Acernese et al. Calibration of advanced Virgo and reconstruction of the detector strain $h(t)$ during the observing run O3. *Class. Quant. Grav.*, 39(4):045006, 2022. doi: 10.1088/1361-6382/ac3c8e.
- [39] F. Acernese et al. Virgo Detector Characterization and Data Quality during the O3 run. *arXiv e-prints*, art. arXiv:2205.01555, May 2022. doi: 10.48550/arXiv.2205.01555.
- [40] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *Class. Quant. Grav.*, 33(17):175012, 2016. doi: 10.1088/0264-9381/33/17/175012.
- [41] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *Astron. Astrophys.*, 594:A13, 2016. doi: 10.1051/0004-6361/201525830.
- [42] Bruce Allen. χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, 2005. doi: 10.1103/PhysRevD.71.062001.
- [43] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85:122006, 2012. doi: 10.1103/PhysRevD.85.122006.

- [44] Pau Amaro-Seoane et al. Laser Interferometer Space Antenna. *arXiv e-prints*, art. arXiv:1702.00786, February 2017. doi: 10.48550/arXiv.1702.00786.
- [45] Fabio Antonini and Mark Gieles. Merger rate of black hole binaries from globular clusters: theoretical error bars and comparison to gravitational wave data from GWTC-2. *Phys. Rev. D*, 102:123016, 2020. doi: 10.1103/PhysRevD.102.123016.
- [46] Fabio Antonini and Frederic A. Rasio. Merging black hole binaries in galactic nuclei: implications for advanced-LIGO detections. *Astrophys. J.*, 831(2):187, 2016. doi: 10.3847/0004-637X/831/2/187.
- [47] K. G. Arun, Alessandra Buonanno, Guillaume Faye, and Evan Ochsner. Higher-order spin effects in the amplitude and phase of gravitational waveforms emitted by inspiraling compact binaries: Ready-to-use gravitational waveforms. *Phys. Rev. D*, 79:104023, 2009. doi: 10.1103/PhysRevD.79.104023. [Erratum: *Phys.Rev.D* 84, 049901 (2011)].
- [48] Gregory Ashton and Colm Talbot. Bilby-MCMC: an MCMC sampler for gravitational-wave inference. *Mon. Not. Roy. Astron. Soc.*, 507(2):2037–2051, 2021. doi: 10.1093/mnras/stab2236.
- [49] Gregory Ashton et al. Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy. *Astrophys. J. Suppl. Ser.*, 241(2):27, 2019. doi: 10.3847/1538-4365/ab06fc.
- [50] F. Aubin et al. The MBTA pipeline for detecting compact binary coalescences in the third LIGO–Virgo observing run. *Class. Quant. Grav.*, 38(9):095004, 2021. doi: 10.1088/1361-6382/abe913.
- [51] Tessa Baker and Mark Trodden. Multimessenger time delays from lensed gravitational waves. *Phys. Rev. D*, 95(6):063512, 2017. doi: 10.1103/PhysRevD.95.063512.
- [52] Matthias Bartelmann and Matteo Maturi. Weak gravitational lensing. *Scholarpedia*, 12(1):32440, January 2017. doi: 10.4249/scholarpedia.32440.
- [53] S. Basak, A. Ganguly, K. Haris, S. Kapadia, A. K. Mehta, and P. Ajith. Constraints on Compact Dark Matter from Gravitational Wave Microlensing. *Astrophys. J.*, 926(2):L28, 2022. doi: 10.3847/2041-8213/ac4dfa.
- [54] K. Belczynski, T. Ryu, R. Perna, E. Berti, T. L. Tanaka, and T. Bulik. On the likelihood of detecting gravitational waves from Population III compact object binaries. *Mon. Not. Roy. Astron. Soc.*, 471(4):4702–4721, 2017. doi: 10.1093/mnras/stx1759.
- [55] K. Belczynski et al. The Effect of Pair-Instability Mass Loss on Black Hole Mergers. *Astron. Astrophys.*, 594:A97, 2016. doi: 10.1051/0004-6361/201628980.

- [56] Krzysztof Belczynski, Vassiliki Kalogera, Frederic A. Rasio, Ronald E. Taam, Andreas Zezas, Tomasz Bulik, Thomas J. Maccarone, and Natalia Ivanova. Compact object modeling with the startrack population synthesis code. *Astrophys. J. Suppl.*, 174:223, 2008. doi: 10.1086/521026.
- [57] Krzysztof Belczynski, Michal Dominik, Tomasz Bulik, Richard O’Shaughnessy, Chris Fryer, and Daniel E. Holz. The effect of metallicity on the detection prospects for gravitational waves. *Astrophys. J. Lett.*, 715:L138, 2010. doi: 10.1088/2041-8205/715/2/L138.
- [58] Krzysztof Belczynski, Daniel E. Holz, Tomasz Bulik, and Richard O’Shaughnessy. The first gravitational-wave source from the isolated evolution of two 40-100 Msun stars. *Nature*, 534:512, 2016. doi: 10.1038/nature18322.
- [59] M. Bernardi, F. Shankar, J. B. Hyde, S. Mei, F. Marulli, and R. K. Sheth. Galaxy luminosities, stellar masses, sizes, velocity dispersions as a function of morphological type. *Mon. Not. Roy. Astron. Soc.*, 404:2087, 2010. doi: 10.1111/j.1365-2966.2010.16425.x.
- [60] T. Blaineau et al. New limits from microlensing on Galactic black holes in the mass range $10 M_{\odot} < M < 1000 M_{\odot}$. *Astron. Astrophys.*, 664:A106, 2022. doi: 10.1051/0004-6361/202243430.
- [61] Luc Blanchet. Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries. *Living Rev. Rel.*, 17:2, 2014. doi: 10.12942/lrr-2014-2.
- [62] Luc Blanchet, Thibault Damour, Bala R. Iyer, Clifford M. Will, and Alan G. Wiseman. Gravitational radiation damping of compact binary systems to second postNewtonian order. *Phys. Rev. Lett.*, 74:3515–3518, 1995. doi: 10.1103/PhysRevLett.74.3515.
- [63] Luc Blanchet, Thibault Damour, Gilles Esposito-Farese, and Bala R. Iyer. Dimensional regularization of the third post-Newtonian gravitational wave generation from two point masses. *Phys. Rev. D*, 71:124004, 2005. doi: 10.1103/PhysRevD.71.124004.
- [64] Roger Blandford and Ramesh Narayan. Fermat’s principle, caustics, and the classification of gravitational lens images. *Astrophys. J.*, 310:568–582, 1986. doi: 10.1086/164709.
- [65] Lumen Boco, A Lapi, S Goswami, F Perrotta, C Baccigalupi, and L Danese. Merging rates of compact binaries in galaxies: perspectives for gravitational wave detections. *The Astrophysical Journal*, 881(2):157, 2019.
- [66] Alejandro Bohé, Sylvain Marsat, and Luc Blanchet. Next-to-next-to-leading order spin–orbit effects in the gravitational wave flux and orbital phasing of

- compact binaries. *Class. Quant. Grav.*, 30:135009, 2013. doi: 10.1088/0264-9381/30/13/135009.
- [67] Alejandro Bohé, Guillaume Faye, Sylvain Marsat, and Edward K. Porter. Quadratic-in-spin effects in the orbital dynamics and gravitational-wave energy flux of compact binaries at the 3PN order. *Class. Quant. Grav.*, 32(19):195010, 2015. doi: 10.1088/0264-9381/32/19/195010.
 - [68] Alejandro Bohé, Mark Hannam, Sascha Husa, Frank Ohme, Michael Puerrer, and Patricia Schmidt. Phenompv2 - technical notes for lal implementation. Technical Report LIGO-T1500602, LIGO Project, 2016. URL <https://dcc.ligo.org/LIGO-T1500602>.
 - [69] Alejandro Bohé et al. Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors. *Phys. Rev. D*, 95(4):044028, 2017. doi: 10.1103/PhysRevD.95.044028.
 - [70] Ian A. Bond et al. OGLE 2003-BLG-235/MOA 2003-BLG-53: A Planetary microlensing event. *Astrophys. J. Lett.*, 606:L155–L158, 2004. doi: 10.1086/420928.
 - [71] J. R. Bond and B. J. Carr. Gravitational waves from a population of binary black holes. *Mon. Not. Roy. Astron. Soc.*, 207:585–609, April 1984. doi: 10.1093/mnras/207.3.585.
 - [72] Yann Bouffanais, Michela Mapelli, Filippo Santoliquido, Nicola Giacobbo, Ugo N. Di Carlo, Sara Rastello, M. Celeste Artale, and Giuliano Iorio. New insights on binary black hole formation channels after GWTC-2: young star clusters versus isolated binaries. *Mon. Not. Roy. Astron. Soc.*, 507(4):5224–5235, 2021. doi: 10.1093/mnras/stab2438.
 - [73] Brendon J. Brewer, Aaron A. Dutton, Tommaso Treu, Matthew W. Auger, Philip J. Marshall, Matteo Barnabe, Adam S. Bolton, David C. Koo, and Leon V. E. Koopmans. The SWELLS survey. III. Disfavouring ‘heavy’ initial mass functions for spiral lens galaxies. *Mon. Not. Roy. Astron. Soc.*, 422:3574, 2012. doi: 10.1111/j.1365-2966.2012.20870.x.
 - [74] T. Broadhurst, J. M. Diego, and G. F. Smoot. A Distant Origin For Magnified LIGO/Virgo Black Holes Implied By Binary Component Masses. *arXiv e-prints*, art. arXiv:2002.08821, February 2020. doi: 10.48550/arXiv.2002.08821.
 - [75] Tom Broadhurst, Jose M. Diego, and III Smoot, George. Reinterpreting Low Frequency LIGO/Virgo Events as Magnified Stellar-Mass Black Holes at Cosmological Distances. *arXiv e-prints*, art. arXiv:1802.05273, February 2018. doi: 10.48550/arXiv.1802.05273.

- [76] Tom Broadhurst, Jose M. Diego, and III Smoot, George F. Twin LIGO/Virgo Detections of a Viable Gravitationally-Lensed Black Hole Merger. *arXiv e-prints*, art. arXiv:1901.03190, January 2019. doi: 10.48550/arXiv.1901.03190.
- [77] Tom Broadhurst, Jose M. Diego, and George F. Smoot. Interpreting LIGO/Virgo “Mass-Gap” events as lensed Neutron Star-Black Hole binaries. *arXiv e-prints*, art. arXiv:2006.13219, June 2020. doi: 10.48550/arXiv.2006.13219.
- [78] Volker Bromm, Paolo S. Coppi, and Richard B. Larson. The formation of the first stars. I. The Primordial star forming cloud. *Astrophys. J.*, 564:23–51, 2002. doi: 10.1086/323947.
- [79] Aaron Buikema et al. Sensitivity and performance of the Advanced LIGO detectors in the third observing run. *Phys. Rev. D*, 102(6):062003, 2020. doi: 10.1103/PhysRevD.102.062003.
- [80] Alessandra Buonanno, Bala Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev. D*, 80:084043, 2009. doi: 10.1103/PhysRevD.80.084043.
- [81] Riccardo Buscicchio, Christopher J. Moore, Geraint Pratten, Patricia Schmidt, Matteo Bianconi, and Alberto Vecchio. Constraining the lensing of binary black holes from their stochastic background. *Phys. Rev. Lett.*, 125(14):141102, 2020. doi: 10.1103/PhysRevLett.125.141102.
- [82] Riccardo Buscicchio, Christopher J. Moore, Geraint Pratten, Patricia Schmidt, and Alberto Vecchio. Constraining the lensing of binary neutron stars from their stochastic background. *Phys. Rev. D*, 102(8):081501, 2020. doi: 10.1103/PhysRevD.102.081501.
- [83] Craig Cahillane et al. Calibration uncertainty for Advanced LIGO’s first and second observing runs. *Phys. Rev. D*, 96(10):102001, 2017. doi: 10.1103/PhysRevD.96.102001.
- [84] Thomas Callister, Maya Fishbach, Daniel Holz, and Will Farr. Shouts and Murmurs: Combining Individual Gravitational-Wave Sources with the Stochastic Background to Measure the History of Binary Black Hole Mergers. *Astrophys. J. Lett.*, 896(2):L32, 2020. doi: 10.3847/2041-8213/ab9743.
- [85] Kipp Cannon, Chad Hanna, and Jacob Peoples. Likelihood-Ratio Ranking Statistic for Compact Binary Coalescence Candidates with Rate Estimation. *arXiv e-prints*, art. arXiv:1504.04632, April 2015.
- [86] Kipp Cannon, Sarah Caudill, Chiwai Chan, Bryce Cousins, Jolien D. E. Creighton, Becca Ewing, Heather Fong, Patrick Godwin, Chad Hanna, Shaun

- Hooper, Rachael Huxford, Ryan Magee, Duncan Meacher, Cody Messick, Soichiro Morisaki, Debnandini Mukherjee, Hiroaki Ohta, Alexander Pace, Stephen Privitera, Iris de Ruiter, Surabhi Sachdev, Leo Singer, Divya Singh, Ron Tapia, Leo Tsukada, Daichi Tsuna, Takuya Tsutsui, Koh Ueno, Aaron Viets, Leslie Wade, and Madeline Wade. GstLAL: A software framework for gravitational wave discovery. *SoftwareX*, 14:100680, June 2021. doi: 10.1016/j.softx.2021.100680.
- [87] Kipp Cannon et al. Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence. *Astrophys. J.*, 748:136, 2012. doi: 10.1088/0004-637X/748/2/136.
 - [88] Shuo Cao, Jingzhao Qi, Zhoujian Cao, Marek Biesiada, Jin Li, Yu Pan, and Zong-Hong Zhu. Direct test of the FLRW metric from strongly lensed gravitational wave observations. *Sci. Rep.*, 9(1):11608, 2019. doi: 10.1038/s41598-019-47616-4.
 - [89] Shuo Cao, Jingzhao Qi, Zhoujian Cao, Marek Biesiada, Wei Cheng, and Zong-Hong Zhu. Direct measurement of the distribution of dark matter with strongly lensed gravitational waves. *Astron. Astrophys.*, 659:L5, 2022. doi: 10.1051/0004-6361/202142694.
 - [90] Zhoujian Cao, Li-Fang Li, and Yan Wang. Gravitational lensing effects on parameter estimation in gravitational wave detection with advanced detectors. *Phys. Rev. D*, 90(6):062003, 2014. doi: 10.1103/PhysRevD.90.062003.
 - [91] Bernard Carr and Florian Kuhnel. Primordial Black Holes as Dark Matter: Recent Developments. *Ann. Rev. Nucl. Part. Sci.*, 70:355–394, 2020. doi: 10.1146/annurev-nucl-050520-125911.
 - [92] Bernard Carr, Kazunori Kohri, Yuuiti Sendouda, and Jun’ichi Yokoyama. Constraints on primordial black holes. *Rept. Prog. Phys.*, 84(11):116902, 2021. doi: 10.1088/1361-6633/ac1e31.
 - [93] Mesut Çalışkan, Neha Anil Kumar, Lingyuan Ji, Jose M. Ezquiaga, Roberto Cotesta, Emanuele Berti, and Marc Kamionkowski. Probing wave-optics effects and low-mass dark matter halos with lensing of gravitational waves from massive black holes. *Phys. Rev. D*, 108(12):123543, 2023. doi: 10.1103/PhysRevD.108.123543.
 - [94] Mesut Çalışkan, Jose María Ezquiaga, Otto A. Hannuksela, and Daniel E. Holz. Lensing or luck? False alarm probabilities for gravitational lensing of gravitational waves. *Phys. Rev. D*, 107(6):063023, 2023. doi: 10.1103/PhysRevD.107.063023.
 - [95] Mesut Çalışkan, Lingyuan Ji, Roberto Cotesta, Emanuele Berti, Marc Kamionkowski, and Sylvain Marsat. Observability of lensing of gravitational waves from massive black hole binaries with LISA. *Phys. Rev. D*, 107(4):043029, 2023. doi: 10.1103/PhysRevD.107.043029.

- [96] Chiwai Chan, Kipp Cannon, Sarah Caudill, Heather Fong, Patrick Godwin, Chad Hanna, Shasvath Kapadia, Ryan Magee, Duncan Meacher, Cody Messick, Siddharth R. Mohite, Soichiro Morisaki, Debnandini Mukherjee, Atsushi Nishizawa, Hiroaki Ohta, Alexander Pace, Surabhi Sachdev, Minori Shikauchi, Leo Singer, Leo Tsukada, Daichi Tsuna, Takuya Tsutsui, and Koh Ueno. Improving the background estimation technique in the GstLAL inspiral pipeline with the time-reversed template bank. *arXiv e-prints*, art. arXiv:2009.03025, September 2020. doi: 10.48550/arXiv.2009.03025.
- [97] Chan, James H. H., Suyu, Sherry H., Sonnenfeld, Alessandro, Jaelani, Anton T., More, Anupreeta, Yonehara, Atsunori, Kubota, Yuriko, Coupon, Jean, Lee, Chien-Hsiu, Oguri, Masamune, Rusu, Cristian E., and Wong, Kenneth C. Survey of gravitationally lensed objects in hsc imaging (sugohi) - iv. lensed quasar search in the hsc survey. *Astronomy & Astrophysics*, 636:A87, 2020. doi: 10.1051/0004-6361/201937030. URL <https://doi.org/10.1051/0004-6361/201937030>.
- [98] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Class. Quant. Grav.*, 21:S1809–S1818, 2004. doi: 10.1088/0264-9381/21/20/024.
- [99] Hsin-Yu Chen, Reed Essick, Salvatore Vitale, Daniel E. Holz, and Erik Katsavounidis. Observational Selection Effects with Ground-based Gravitational Wave Detectors. *Astrophys. J.*, 835(1):31, 2017. doi: 10.3847/1538-4357/835/1/31.
- [100] Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. *arXiv e-prints*, art. arXiv:1603.02754, March 2016. doi: 10.48550/arXiv.1603.02754.
- [101] Mark H. Y. Cheung, Joseph Gais, Otto A. Hannuksela, and Tjonnie G. F. Li. Stellar-mass microlensing of gravitational waves. *Mon. Not. Roy. Astron. Soc.*, 503(3):3326–3336, 2021. doi: 10.1093/mnras/stab579.
- [102] Yun-Young Choi, Changbom Park, and Michael S. Vogeley. Internal and Collective Properties of Galaxies in the Sloan Digital Sky Survey. *Astrophys. J.*, 658:884–897, 2007. doi: 10.1086/511060.
- [103] Yun-Young Choi, Changbom Park, and Michael S. Vogeley. Internal and Collective Properties of Galaxies in the Sloan Digital Sky Survey. *Astrophys. J.*, 658(2):884–897, April 2007. doi: 10.1086/511060.
- [104] Nelson Christensen and Renate Meyer. Parameter estimation with gravitational waves. *Rev. Mod. Phys.*, 94(2):025001, 2022. doi: 10.1103/RevModPhys.94.025001.
- [105] Pierre Christian, Salvatore Vitale, and Abraham Loeb. Detecting Stellar Lensing of Gravitational Waves with Ground-Based Observatories. *Phys. Rev. D*, 98(10):103022, 2018. doi: 10.1103/PhysRevD.98.103022.

- [106] Qi Chu et al. SPIIR online coherent pipeline to search for gravitational waves from compact binary coalescences. *Phys. Rev. D*, 105(2):024023, 2022. doi: 10.1103/PhysRevD.105.024023.
- [107] Douglas Clowe, Anthony Gonzalez, and Maxim Markevitch. Weak lensing mass reconstruction of the interacting cluster 1E0657-558: Direct evidence for the existence of dark matter. *Astrophys. J.*, 604:596–603, 2004. doi: 10.1086/381970.
- [108] Dan Coe et al. CLASH: Three Strongly Lensed Images of a Candidate $z \sim 11$ Galaxy. *Astrophys. J.*, 762:32, 2013. doi: 10.1088/0004-637X/762/1/32.
- [109] Marta Colleoni, Maite Mateu-Lucena, Héctor Estellés, Cecilio García-Quirós, David Keitel, Geraint Pratten, Antoni Ramos-Buades, and Sascha Husa. Towards the routine use of subdominant harmonics in gravitational-wave inference: Reanalysis of GW190412 with generation X waveform models. *Phys. Rev. D*, 103(2):024029, 2021. doi: 10.1103/PhysRevD.103.024029.
- [110] Thomas E. Collett and David Bacon. Testing the speed of gravitational waves over cosmological distances with strong gravitational lensing. *Phys. Rev. Lett.*, 118(9):091101, 2017. doi: 10.1103/PhysRevLett.118.091101.
- [111] A.B. Congdon and C.R. Keeton. *Principles of Gravitational Lensing: Light Deflection as a Probe of Astrophysics and Cosmology*. Springer Praxis Books. Springer International Publishing, 2018. ISBN 9783030021221. URL <https://books.google.com/books?id=kt58DwAAQBAJ>.
- [112] Liam Connor and Vikram Ravi. Stellar prospects for FRB gravitational lensing. *Mon. Not. Roy. Astron. Soc.*, 521(3):4024–4038, 2023. doi: 10.1093/mnras/stad667.
- [113] Neil Cornish, Laura Sampson, Nicolas Yunes, and Frans Pretorius. Gravitational Wave Tests of General Relativity with the Parameterized Post-Einsteinian Framework. *Phys. Rev. D*, 84:062003, 2011. doi: 10.1103/PhysRevD.84.062003.
- [114] Neil J. Cornish and Tyson B. Littenberg. BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches. *Class. Quant. Grav.*, 32(13):135012, 2015. doi: 10.1088/0264-9381/32/13/135012.
- [115] Neil J. Cornish, Tyson B. Littenberg, Bence Bécsy, Katerina Chatziioannou, James A. Clark, Sudarshan Ghonge, and Margaret Millhouse. BayesWave analysis pipeline in the era of gravitational wave observations. *Phys. Rev. D*, 103(4):044006, 2021. doi: 10.1103/PhysRevD.103.044006.
- [116] Paolo Cremonese, Jose María Ezquiaga, and Vincenzo Salzano. Breaking the mass-sheet degeneracy with gravitational wave interference in lensed events. *Phys. Rev. D*, 104(2):023503, 2021. doi: 10.1103/PhysRevD.104.023503.

- [117] Pedro V. P. Cunha and Carlos A. R. Herdeiro. Shadows and strong gravitational lensing: a brief review. *Gen. Rel. Grav.*, 50(4):42, 2018. doi: 10.1007/s10714-018-2361-9.
- [118] Curt Cutler and Eanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral wave form? *Phys. Rev. D*, 49:2658–2697, 1994. doi: 10.1103/PhysRevD.49.2658.
- [119] Liang Dai and Tejaswi Venumadhav. On the waveforms of gravitationally lensed gravitational waves. *arXiv e-prints*, art. arXiv:1702.04724, February 2017. doi: 10.48550/arXiv.1702.04724.
- [120] Liang Dai, Tejaswi Venumadhav, and Kris Sigurdson. Effect of lensing magnification on the apparent distribution of black hole mergers. *Phys. Rev. D*, 95(4):044011, 2017. doi: 10.1103/PhysRevD.95.044011.
- [121] Liang Dai, Shun-Sheng Li, Barak Zackay, Shude Mao, and Youjun Lu. Detecting Lensing-Induced Diffraction in Astrophysical Gravitational Waves. *Phys. Rev. D*, 98(10):104029, 2018. doi: 10.1103/PhysRevD.98.104029.
- [122] Liang Dai, Barak Zackay, Tejaswi Venumadhav, Javier Roulet, and Matias Zaldarriaga. Search for Lensed Gravitational Waves Including Morse Phase Information: An Intriguing Candidate in O2. *arXiv e-prints*, art. arXiv:2007.12709, July 2020. doi: 10.48550/arXiv.2007.12709.
- [123] Tito Dal Canton et al. Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors. *Phys. Rev. D*, 90(8):082004, 2014. doi: 10.1103/PhysRevD.90.082004.
- [124] Charles Dalang, Giulia Cusin, and Macarena Lagos. Polarization distortions of lensed gravitational waves. *Phys. Rev. D*, 105(2):024005, 2022. doi: 10.1103/PhysRevD.105.024005.
- [125] Thibault Damour and Alessandro Nagar. The Effective-One-Body Approach to the General Relativistic Two Body Problem. *Lect. Notes Phys.*, 905:273–312, 2016. doi: 10.1007/978-3-319-19416-5_7.
- [126] Thibault Damour, Piotr Jaranowski, and Gerhard Schaefel. On the determination of the last stable orbit for circular general relativistic binaries at the third postNewtonian approximation. *Phys. Rev. D*, 62:084011, 2000. doi: 10.1103/PhysRevD.62.084011.
- [127] Gareth S. Davies, Thomas Dent, Márton Tápai, Ian Harry, Connor McIsaac, and Alexander H. Nitz. Extending the PyCBC search for gravitational waves from compact binary mergers to a global network. *Phys. Rev. D*, 102(2):022004, 2020. doi: 10.1103/PhysRevD.102.022004.

- [128] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall. Improving the Sensitivity of Advanced LIGO Using Noise Subtraction. *Class. Quant. Grav.*, 36(5):055011, 2019. doi: 10.1088/1361-6382/ab01c5.
- [129] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton. Subtracting glitches from gravitational-wave detector data during the third LIGO-Virgo observing run. *Class. Quant. Grav.*, 39(24):245013, 2022. doi: 10.1088/1361-6382/aca238.
- [130] Derek Davis and Marissa Walker. Detector Characterization and Mitigation of Noise in Ground-Based Gravitational-Wave Interferometers. *Galaxies*, 10(1):12, 2022. doi: 10.3390/galaxies10010012.
- [131] Derek Davis et al. LIGO detector characterization in the second and third observing runs. *Class. Quant. Grav.*, 38(13):135014, 2021. doi: 10.1088/1361-6382/abfd85.
- [132] V De Luca, G Franciolini, P Pani, and A Riotto. The evolution of primordial black holes and their final observable spins. *Journal of Cosmology and Astroparticle Physics*, 2020(04):052, 2020.
- [133] Francesco De Paolis, Mosè Giordano, Gabriele Ingrosso, Luigi Manni, Achille Nucita, and Francesco Strafella. The scales of gravitational lensing. *Universe*, 2(1), 2016. ISSN 2218-1997. doi: 10.3390/universe2010006. URL <https://www.mdpi.com/2218-1997/2/1/6>.
- [134] Shuji Deguchi and William D. Watson. Wave effects in gravitational lensing of electromagnetic radiation. *Phys. Rev. D*, 34:1708–1718, 1986. doi: 10.1103/PhysRevD.34.1708.
- [135] Walter Del Pozzo, Katherine Grover, Ilya Mandel, and Alberto Vecchio. Testing general relativity with compact coalescing binaries: comparing exact and predictive methods to compute the Bayes factor. *Class. Quant. Grav.*, 31(20):205006, 2014. doi: 10.1088/0264-9381/31/20/205006.
- [136] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.
- [137] Thomas Dent. "extending the pycbc pastro calculation to a global network", 2021. URL https://dcc.ligo.org/public/0173/T2100060/001/pycbc_multi_p_astro.pdf.
- [138] J.M. Diego, O.A. Hannuksela, P.L. Kelly, T. Broadhurst, K. Kim, T.G.F. Li, G.F. Smoot, and G. Pagano. Observational signatures of microlensing in gravitational waves at LIGO/Virgo frequencies. *Astron. Astrophys.*, 627:A130, 2019. doi: 10.1051/0004-6361/201935490.

- [139] Jose M. Diego. Constraining the abundance of primordial black holes with gravitational lensing of gravitational waves at LIGO frequencies. *Phys. Rev. D*, 101(12):123512, 2020. doi: 10.1103/PhysRevD.101.123512.
- [140] Tim Dietrich et al. Matter imprints in waveform models for neutron star binaries: Tidal and self-spin effects. *Phys. Rev. D*, 99(2):024029, 2019. doi: 10.1103/PhysRevD.99.024029.
- [141] Scott Dodelson. *Gravitational Lensing*. Cambridge University Press, 2017. doi: 10.1017/9781316424254.
- [142] V. I. Dokuchaev and N. O. Nazarova. Gravitational Lensing of a star by a rotating black hole. *JETP Lett.*, 106(10):637–642, 2017. doi: 10.1134/S0021364017220088. [Pisma Zh. Eksp. Teor. Fiz.106,no.10,609(2017)].
- [143] Michal Dominik, Krzysztof Belczynski, Christopher Fryer, Daniel E. Holz, Emanuele Berti, Tomasz Bulik, Ilya Mandel, and Richard O’Shaughnessy. Double Compact Objects II: Cosmological Merger Rates. *Astrophys. J.*, 779: 72, 2013. doi: 10.1088/0004-637X/779/1/72.
- [144] Marco Drago, Sergey Klimenko, Claudia Lazzaro, Edoardo Milotti, Guenakh Mitselmakher, Valentin Necula, Brendan O’Brian, Giovanni Andrea Prodi, Francesco Salemi, Marek Szczepanczyk, Shubhanshu Tiwari, Vaibhav Tiwari, V. Gayathri, Gabriele Vedovato, and Igor Yakushin. coherent WaveBurst, a pipeline for unmodeled gravitational-wave data analysis. *SoftwareX*, 14:100678, June 2021. doi: 10.1016/j.softx.2021.100678.
- [145] J. C. Driggers et al. Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO. *Phys. Rev. D*, 99(4):042001, 2019. doi: 10.1103/PhysRevD.99.042001.
- [146] J. Ehlers and P. Schneider. Gravitational lensing. In *13th Conference on General Relativity and Gravitation (GR-13)*, pages 21–40, 1993.
- [147] Albert Einstein. Lens-like action of a star by the deviation of light in the gravitational field. *Science*, 84(2188):506–507, 1936. doi: 10.1126/science.84.2188.506. URL <https://www.science.org/doi/abs/10.1126/science.84.2188.506>.
- [148] J. J. Eldridge, E. R. Stanway, and Petra N. Tang. A consistent estimate for Gravitational Wave and Electromagnetic transient rates. *Mon. Not. Roy. Astron. Soc.*, 482(1):870–880, 2019. doi: 10.1093/mnras/sty2714.
- [149] Reed Essick, Patrick Godwin, Chad Hanna, Lindy Blackburn, and Erik Katsavounidis. iDQ: Statistical Inference of Non-Gaussian Noise with Auxiliary Degrees of Freedom in Gravitational-Wave Detectors. *arXiv e-prints*, art. arXiv:2005.12761, May 2020. doi: 10.48550/arXiv.2005.12761.

- [150] A. Esteban-Gutiérrez, E. Mediavilla, J. Jiménez-Vicente, N. Agües-Paszkowsky, J. A. Muñoz, and S. Heydenreich. Abundance of LIGO/Virgo Black Holes from Microlensing Observations of Quasars with Reverberation Mapping Size Estimates. *Astrophys. J.*, 929(2):123, 2022. doi: 10.3847/1538-4357/ac57c5.
- [151] Héctor Estellés, Marta Colleoni, Cecilio García-Quirós, Sascha Husa, David Keitel, Maite Mateu-Lucena, Maria de Lluc Planas, and Antoni Ramos-Buades. New twists in compact binary waveform modeling: A fast time-domain model for precession. *Phys. Rev. D*, 105(8):084040, 2022. doi: 10.1103/PhysRevD.105.084040.
- [152] Héctor Estellés et al. A Detailed Analysis of GW190521 with Phenomenological Waveform Models. *Astrophys. J.*, 924(2):79, 2022. doi: 10.3847/1538-4357/ac33a0.
- [153] Becca Ewing et al. Performance of the low-latency GstLAL inspiral search towards LIGO, Virgo, and KAGRA’s fourth observing run. *Phys. Rev. D*, 109(4):042008, 2024. doi: 10.1103/PhysRevD.109.042008.
- [154] J. M. Ezquiaga, Daniel E. Holz, Wayne Hu, Macarena Lagos, and Robert M. Wald. Phase effects from strong gravitational lensing of gravitational waves. *Phys. Rev. D*, 103:6, 2021. doi: 10.1103/PhysRevD.103.064047.
- [155] Jose María Ezquiaga and Miguel Zumalacárregui. Gravitational wave lensing beyond general relativity: birefringence, echoes and shadows. *Phys. Rev. D*, 102(12):124048, 2020. doi: 10.1103/PhysRevD.102.124048.
- [156] Jose Maria Ezquiaga, Wayne Hu, Macarena Lagos, Meng-Xiang Lin, and Fei Xu. Modified gravitational wave propagation with higher modes and its degeneracies with lensing. *Journal of Cosmology and Astroparticle Physics*, 2022(08):016, 2022.
- [157] Jose María Ezquiaga, Wayne Hu, and Rico K. L. Lo. Identifying strongly lensed gravitational waves through their phase consistency. *Phys. Rev. D*, 108(10):103520, 2023. doi: 10.1103/PhysRevD.108.103520.
- [158] Malcolm Fairbairn, Juan Urrutia, and Ville Vaskonen. Microlensing of gravitational waves by dark matter structures. *Journal of Cosmology and Astroparticle Physics*, 2023(07):007, 2023.
- [159] Will M. Farr, Jonathan R. Gair, Ilya Mandel, and Curt Cutler. Counting And Confusion: Bayesian Rate Estimation With Multiple Populations. *Phys. Rev. D*, 91(2):023005, 2015. doi: 10.1103/PhysRevD.91.023005.
- [160] Nicholas Farrow, Xing-Jiang Zhu, and Eric Thrane. The mass distribution of Galactic double neutron stars. *Astrophys. J.*, 876(1):18, 2019. doi: 10.3847/1538-4357/ab12e3.

- [161] Guillaume Faye, Sylvain Marsat, Luc Blanchet, and Bala R. Iyer. The third and a half post-Newtonian gravitational wave quadrupole mode for quasi-circular inspiralling compact binaries. *Class. Quant. Grav.*, 29:175004, 2012. doi: 10.1088/0264-9381/29/17/175004.
- [162] Lee Samuel Finn and David F. Chernoff. Observing binary inspiral in gravitational radiation: One interferometer. *Phys. Rev. D*, 47:2198–2219, 1993. doi: 10.1103/PhysRevD.47.2198.
- [163] Irene Fiori et al. The Hunt for Environmental Noise in Virgo during the Third Observing Run. *Galaxies*, 8(4):82, 2020. doi: 10.3390/galaxies8040082.
- [164] Maya Fishbach and Daniel E. Holz. Where Are LIGO’s Big Black Holes? *Astrophys. J. Lett.*, 851(2):L25, 2017. doi: 10.3847/2041-8213/aa9bf6.
- [165] C. Fletcher et al. A Joint Fermi-GBM and Swift-BAT Analysis of Gravitational-wave Candidates from the Third Gravitational-wave Observing Run. *Astrophys. J.*, 964(2):149, 2024. doi: 10.3847/1538-4357/ad1eed.
- [166] Heather Kin Yee Fong. *From Simulations to Signals: Analyzing Gravitational Waves from Compact Binary Coalescences*. PhD thesis, Department of Physics, University of Toronto, 7 2018.
- [167] Giacomo Fragione and Joseph Silk. Repeated mergers and ejection of black holes within nuclear star clusters. *Mon. Not. Roy. Astron. Soc.*, 498(4):4591–4604, 2020. doi: 10.1093/mnras/staa2629.
- [168] Masataka Fukugita and Edwin L. Turner. Gravitational lensing frequencies: Galaxy cross-sections and selection effects. *Mon. Not. Roy. Astron. Soc.*, 253: 99–106, 1991.
- [169] Juan Garcia-Bellido, Sebastien Clesse, and Pierre Fleury. Primordial black holes survive SN lensing constraints. *Phys. Dark Univ.*, 20:95–100, 2018. doi: 10.1016/j.dark.2018.04.005.
- [170] Cecilio García-Quirós, Marta Colleoni, Sascha Husa, Héctor Estellés, Geraint Pratten, Antoni Ramos-Buades, Maite Mateu-Lucena, and Rafel Jaume. Multimode frequency-domain model for the gravitational wave signal from non-precessing black-hole binaries. *Phys. Rev. D*, 102(6):064002, 2020. doi: 10.1103/PhysRevD.102.064002.
- [171] Ángel Garrón and David Keitel. Waveform systematics in identifying strongly gravitationally lensed gravitational waves: posterior overlap method. *Class. Quant. Grav.*, 41(1):015005, 2024. doi: 10.1088/1361-6382/ad0b9b.
- [172] Raphael Gavazzi, Tommaso Treu, Philip J Marshall, Florence Brault, and Andrea Ruff. The sl2s galaxy-scale gravitational lens sample. i. the alignment of mass and light in massive early-type galaxies at $z=0.2\text{--}0.9$. *The Astrophysical Journal*, 761(2):170, 2012.

- [173] Gerard Gilmore and Gudrun Tausch-Pebody. The 1919 eclipse results that verified general relativity and their later detractors: a story re-told. *Notes and Records: the Royal Society Journal of the History of Science*, 76(1):155–180, 2022. doi: 10.1098/rsnr.2020.0040. URL <https://royalsocietypublishing.org/doi/abs/10.1098/rsnr.2020.0040>.
- [174] Jean-Marc Ginoux. Albert einstein and the doubling of the deflection of light. *Foundations of Science*, 27:1–22, 02 2021. doi: 10.1007/s10699-021-09783-4.
- [175] Given. 冬のはなし, 2020.
- [176] Patrick Godwin et al. Incorporation of Statistical Data Quality Information into the GstLAL Search Analysis. 10 2020.
- [177] Srashti Goyal, Harikrishnan D., Shasvath J. Kapadia, and Parameswaran Ajith. Rapid identification of strongly lensed gravitational-wave events with machine learning. *Phys. Rev. D*, 104(12):124057, 2021. doi: 10.1103/PhysRevD.104.124057.
- [178] Srashti Goyal, K. Haris, Ajit Kumar Mehta, and Parameswaran Ajith. Testing the nature of gravitational-wave polarizations using strongly lensed signals. *Phys. Rev. D*, 103(2):024038, 2021. doi: 10.1103/PhysRevD.103.024038.
- [179] Srashti Goyal, Aditya Vijaykumar, Jose Maria Ezquiaga, and Miguel Zumalacarregui. Probing lens-induced gravitational-wave birefringence as a test of general relativity. *Phys. Rev. D*, 108(2):024052, 2023. doi: 10.1103/PhysRevD.108.024052.
- [180] Srashti Goyal, Shasvath J. Kapadia, Jean-Rene Cudell, Alvin K. Y. Li, and Juno C. L. Chan. Rapid method for preliminary identification of subthreshold strongly lensed counterparts to superthreshold gravitational-wave events. *Phys. Rev. D*, 109(2):023028, 2024. doi: 10.1103/PhysRevD.109.023028.
- [181] GWOSC. GWTC-2 Data Release, 2020.
- [182] GWOSC. GWTC-3 Data Release, 2021.
- [183] Chad Hanna et al. Fast evaluation of multidetector consistency for real-time gravitational wave searches. *Phys. Rev. D*, 101(2):022003, 2020. doi: 10.1103/PhysRevD.101.022003.
- [184] Mark Hannam, Patricia Schmidt, Alejandro Bohé, Leïla Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Pürrer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014. doi: 10.1103/PhysRevLett.113.151101.
- [185] Mark Hannam et al. General-relativistic precession in a black-hole binary. *Nature*, 610(7933):652–655, 2022. doi: 10.1038/s41586-022-05212-z.

- [186] O. A. Hannuksela, K. Haris, K. K. Y. Ng, S. Kumar, A. K. Mehta, D. Keitel, T. G. F. Li, and P. Ajith. Search for gravitational lensing signatures in LIGO-Virgo binary black hole events. *Astrophys. J. Lett.*, 874(1):L2, 2019. doi: 10.3847/2041-8213/ab0c0f.
- [187] Otto A. Hannuksela, Thomas E. Collett, Mesut Çalışkan, and Tjonnie G. F. Li. Localizing merging black holes with sub-arcsecond precision using gravitational-wave lensing. *Mon. Not. Roy. Astron. Soc.*, 498(3):3395–3402, 2020. doi: 10.1093/mnras/staa2577.
- [188] K. Haris, Ajit Kumar Mehta, Sumit Kumar, Tejaswi Venumadhav, and Parameswaran Ajith. Identifying strongly lensed gravitational wave signals from binary black hole mergers. *arXiv e-prints*, art. arXiv:1807.07062, July 2018. doi: 10.48550/arXiv.1807.07062.
- [189] Carl W. Helstrom and Fred L. Wilson. Statistical Theory Of Signal Detection. *Physics Today*, 23(6):73, January 1970. doi: 10.1063/1.3022201.
- [190] Wen-Rui Hu and Yue-Liang Wu. The Taiji Program in Space for gravitational wave physics and the nature of gravity. *Natl. Sci. Rev.*, 4(5):685–686, 2017. doi: 10.1093/nsr/nwx116.
- [191] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. *arXiv e-prints*, art. arXiv:1608.06993, August 2016.
- [192] K. Hurley, A. E. Tsvetkova, D. S. Svinkin, R. L. Aptekar, D. D. Frederiks, S. V. Golenetskii, A. A. Kokomov, A. V. Kozlova, A. L. Lysenko, M. V. Ulanov, T. L. Cline, I. G. Mitrofanov, D. Golovin, M. L. Litvak, A. B. Sanin, W. Boynton, K. Harshman, C. Fellows, R. Starr, A. Rau, A. von Kienlin, and X. Zhang. A search for gravitationally lensed gamma-ray bursts in the data of the interplanetary network and konus-wind. *The Astrophysical Journal*, 871(1):121, jan 2019. doi: 10.3847/1538-4357/aaf645. URL <https://dx.doi.org/10.3847/1538-4357/aaf645>.
- [193] Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev. D*, 93(4):044006, 2016. doi: 10.1103/PhysRevD.93.044006.
- [194] Miho N. Ishigaki, Nozomu Tominaga, Chiaki Kobayashi, and Ken’ichi Nomoto. The Initial Mass Function of the First Stars Inferred from Extremely Metal-poor Stars. *Astrophys. J.*, 857(1):46, April 2018. doi: 10.3847/1538-4357/aab3de.
- [195] Anton T Jaelani, Anupreeta More, Masamune Oguri, Alessandro Sonnenfeld, Sherry H Suyu, Cristian E Rusu, Kenneth C Wong, James H H Chan,

- Issha Kayo, Chien-Hsiu Lee, Dani C-Y Chao, Jean Coupon, Kaiki T Inoue, and Toshifumi Futamase. Survey of Gravitationally lensed Objects in HSC Imaging (SuGOHI) – V. Group-to-cluster scale lens search from the HSC–SSP Survey. *Monthly Notices of the Royal Astronomical Society*, 495(1):1291–1310, 04 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1062. URL <https://doi.org/10.1093/mnras/staa1062>.
- [196] Anton T Jaelani, Cristian E Rusu, Issha Kayo, Anupreeta More, Alessandro Sonnenfeld, John D Silverman, Malte Schramm, Timo Anguita, Naohisa Inada, Daichi Kondo, Paul L Schechter, Khee-Gan Lee, Masamune Oguri, James H H Chan, Kenneth C Wong, and Kaiki T Inoue. Survey of Gravitationally Lensed Objects in HSC Imaging (SuGOHI) – VII. Discovery and confirmation of three strongly lensed quasars†. *Monthly Notices of the Royal Astronomical Society*, 502(1):1487–1493, 01 2021. ISSN 0035-8711. doi: 10.1093/mnras/stab145. URL <https://doi.org/10.1093/mnras/stab145>.
- [197] Souvik Jana, Shasvath J. Kapadia, Tejaswi Venumadhav, and Parameswaran Ajith. Cosmography Using Strongly Lensed Gravitational Waves from Binary Black Holes. *Phys. Rev. Lett.*, 130(26):261401, 2023. doi: 10.1103/PhysRevLett.130.261401.
- [198] Justin Janquart, Otto A. Hannuksela, Haris K., and Chris Van Den Broeck. A fast and precise methodology to search for and analyse strongly lensed gravitational-wave events. *Mon. Not. Roy. Astron. Soc.*, 506(4):5430–5438, 2021. doi: 10.1093/mnras/stab1991.
- [199] Justin Janquart, Eungwang Seo, Otto A. Hannuksela, Tjonnie G. F. Li, and Chris Van Den Broeck. On the Identification of Individual Gravitational-wave Image Types of a Lensed System Using Higher-order Modes. *Astrophys. J. Lett.*, 923(1):L1, 2021. doi: 10.3847/2041-8213/ac3bcf.
- [200] Justin Janquart, Otto A. Hannuksela, K. Haris, and Chris Van den Broeck. GOLUM: A fast and precise methodology to search for, and analyze, strongly lensed gravitational-wave events. In *56th Rencontres de Moriond on Gravitation*, 3 2022.
- [201] Justin Janquart, Anupreeta More, and Chris Van Den Broeck. Ordering the confusion: a study of the impact of lens models on gravitational-wave strong lensing detection capabilities. *Mon. Not. Roy. Astron. Soc.*, 519(2): 2046–2059, 2022. doi: 10.1093/mnras/stac3660.
- [202] Justin Janquart, K. Haris, Otto A. Hannuksela, and Chris Van Den Broeck. The return of GOLUM: improving distributed joint parameter estimation for strongly lensed gravitational waves. *Mon. Not. Roy. Astron. Soc.*, 526(2): 3088–3098, 2023. doi: 10.1093/mnras/stad2838.

- [203] Justin Janquart et al. Follow-up analyses to the O3 LIGO–Virgo–KAGRA lensing searches. *Mon. Not. Roy. Astron. Soc.*, 526(3):3832–3860, 2023. doi: 10.1093/mnras/stad2909.
- [204] David E. Johnston et al. SDSS J0903+5028: A New gravitational lens. *Astron. J.*, 126:2281, 2003. doi: 10.1086/379001.
- [205] Sunghoon Jung and Chang Sub Shin. Gravitational-Wave Fringes at LIGO: Detecting Compact Dark Matter by Gravitational Lensing. *Phys. Rev. Lett.*, 122(4):041103, 2019. doi: 10.1103/PhysRevLett.122.041103.
- [206] Shasvath J. Kapadia et al. A self-consistent method to estimate the rate of compact binary coalescences with a Poisson mixture model. *Class. Quant. Grav.*, 37(4):045007, 2020. doi: 10.1088/1361-6382/ab5f2d.
- [207] S. Karki et al. The Advanced LIGO Photon Calibrators. *Rev. Sci. Instrum.*, 87(11):114503, 2016. doi: 10.1063/1.4967303.
- [208] P. Kelly. Supernova Refsdal: Time Delay Measurements and H0 Constraints. In *American Astronomical Society Meeting Abstracts*, volume 53 of *American Astronomical Society Meeting Abstracts*, page 410.05, January 2021.
- [209] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev. D*, 93(4):044007, 2016. doi: 10.1103/PhysRevD.93.044007.
- [210] Sebastian Khan, Frank Ohme, Katerina Chatziioannou, and Mark Hannam. Including higher order multipoles in gravitational-wave models for precessing binary black holes. *Phys. Rev. D*, 101(2):024056, 2020. doi: 10.1103/PhysRevD.101.024056.
- [211] Chase Kimball et al. Evidence for Hierarchical Black Hole Mergers in the Second LIGO–Virgo Gravitational Wave Catalog. *Astrophys. J. Lett.*, 915(2):L35, 2021. doi: 10.3847/2041-8213/ac0aef.
- [212] S. Klimenko and Guenakh Mitselmakher. A wavelet method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 21:S1819–S1830, 2004. doi: 10.1088/0264-9381/21/20/025.
- [213] S. Klimenko, S. Mohanty, Malik Rakhmanov, and Guenakh Mitselmakher. Constraint likelihood analysis for a network of gravitational wave detectors. *Phys. Rev. D*, 72:122002, 2005. doi: 10.1103/PhysRevD.72.122002.
- [214] S. Klimenko, S. Mohanty, Malik Rakhmanov, and Guenakh Mitselmakher. Constraint likelihood method: Generalization for colored noise. *J. Phys. Conf. Ser.*, 32:12–17, 2006. doi: 10.1088/1742-6596/32/1/003.

- [215] S. Klimenko, I. Yakushin, A. Mercer, and Guenakh Mitselmakher. Coherent method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 25: 114029, 2008. doi: 10.1088/0264-9381/25/11/114029.
- [216] S. Klimenko, G. Vedovato, M. Drago, G. Mazzolo, G. Mitselmakher, C. Pankow, G. Prodi, V. Re, F. Salemi, and I. Yakushin. Localization of gravitational wave sources with networks of advanced detectors. *Phys. Rev. D*, 83:102001, 2011. doi: 10.1103/PhysRevD.83.102001.
- [217] S. Klimenko et al. Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. *Phys. Rev. D*, 93(4): 042004, 2016. doi: 10.1103/PhysRevD.93.042004.
- [218] LVE Koopmans, A Bolton, T Treu, O Czoske, MW Auger, M Barnabe, S Vegetti, R Gavazzi, LA Moustakas, and S Burles. The structure and dynamics of massive early-type galaxies: On homology, isothermality, and isotropy inside one effective radius. *The Astrophysical Journal*, 703(1):L51, 2009.
- [219] I Kowalska, T Bulik, and K Belczynski. Gravitational wave background from population iii binaries. *Astronomy & Astrophysics*, 541:A120, 2012.
- [220] Kwun-Hang Lai, Otto A. Hannuksela, Antonio Herrera-Martín, Jose M. Diego, Tom Broadhurst, and Tjonnie G.F. Li. Discovering intermediate-mass black hole lenses through gravitational wave lensing. *Phys. Rev. D*, 98(8): 083005, 2018. doi: 10.1103/PhysRevD.98.083005.
- [221] J. Lange et al. Parameter estimation method that directly compares gravitational wave observations to numerical relativity. *Phys. Rev. D*, 96(10): 104041, 2017. doi: 10.1103/PhysRevD.96.104041.
- [222] Alvin K. Y. Li, Juno C. L. Chan, Heather Fong, Aidan H. Y. Chong, Alan J. Weinstein, and Jose M. Ezquiaga. TESLA-X: An effective method to search for sub-threshold lensed gravitational waves with a targeted population model. *arXiv e-prints*, art. arXiv:2311.06416, November 2023. doi: 10.48550/arXiv.2311.06416.
- [223] Alvin K. Y. Li, Rico K. L. Lo, Surabhi Sachdev, Juno C. L. Chan, E. T. Lin, Tjonnie G. F. Li, and Alan J. Weinstein. Targeted subthreshold search for strongly lensed gravitational-wave events. *Phys. Rev. D*, 107(12):123014, 2023. doi: 10.1103/PhysRevD.107.123014.
- [224] Shun-Sheng Li, Shude Mao, Yuetong Zhao, and Youjun Lu. Gravitational lensing of gravitational waves: A statistical perspective. *Mon. Not. Roy. Astron. Soc.*, 476(2):2220–2229, 2018. doi: 10.1093/mnras/sty411.
- [225] Yufeng Li, Xilong Fan, and Lijun Gou. Constraining Cosmological Parameters in the FLRW Metric with Lensed GW+EM Signals. *Astrophys. J.*, 873 (1):37, 2019. doi: 10.3847/1538-4357/ab037e.

- [226] Zheng-Xiang Li, He Gao, Xu-Heng Ding, Guo-Jian Wang, and Bing Zhang. Strongly lensed repeating fast radio bursts as precision probes of the universe. *Nature Commun.*, 9(1):3833, 2018. doi: 10.1038/s41467-018-06303-0.
- [227] Kai Liao, Xi-Long Fan, Xu-Heng Ding, Marek Biesiada, and Zong-Hong Zhu. Precision cosmology from future lensed gravitational wave and electromagnetic signals. *Nature Commun.*, 8(1):1148, 2017. doi: 10.1038/s41467-017-01152-9. [Erratum: *Nature Commun.* 8, 2136 (2017)].
- [228] LIGO Scientific Collaboration and Virgo Collaboration. Data release for "search for lensing signatures in the gravitational-wave observations from the first half of ligo-virgo's third observing run". <https://dcc.ligo.org/P2100173/public>, 2021.
- [229] Shi-Jie Lin et al. A Search for Millilensing Gamma-Ray Bursts in the Observations of Fermi GBM. *Astrophys. J.*, 931(1):4, 2022. doi: 10.3847/1538-4357/ac6505.
- [230] Anna Liu, Isaac C. F. Wong, Samson H. W. Leong, Anupreeta More, Otto A. Hannuksela, and Tjonnie G. F. Li. Exploring the hidden Universe: a novel phenomenological approach for recovering arbitrary gravitational-wave millilensing configurations. *Mon. Not. Roy. Astron. Soc.*, 525(3):4149–4160, 2023. doi: 10.1093/mnras/stad1302.
- [231] Bin Liu, Zhengxiang Li, and Zong-Hong Zhu. Complementary constraints on dark energy equation of state from strongly lensed gravitational wave. *Mon. Not. Roy. Astron. Soc.*, 487(2):1980–1985, 2019. doi: 10.1093/mnras/stz1179.
- [232] Boyuan Liu and Volker Bromm. Gravitational waves from Population III binary black holes formed by dynamical capture. *Mon. Not. Roy. Astron. Soc.*, 495(2):2475–2495, 2020. doi: 10.1093/mnras/staa1362.
- [233] Xiaoshu Liu, Ignacio Magana Hernandez, and Jolien Creighton. Identifying strong gravitational-wave lensing during the second observing run of Advanced LIGO and Advanced Virgo. *Astrophys. J.*, 908(1):97, 2021. doi: 10.3847/1538-4357/abd7eb.
- [234] Rico K. L. Lo. denmarf: a Python package for density estimation using masked autoregressive flow. *arXiv e-prints*, art. arXiv:2305.14379, May 2023. doi: 10.48550/arXiv.2305.14379.
- [235] Rico K. L. Lo and Ignacio Magana Hernandez. Bayesian statistical framework for identifying strongly lensed gravitational-wave signals. *Phys. Rev. D*, 107(12):123015, 2023. doi: 10.1103/PhysRevD.107.123015.
- [236] Rico K. L. Lo and Masamune Oguri. Incorporating Astrophysical Information in the Identification of Strongly-Lensed Gravitational Waves. in prep.

- [237] Jing Luan, Shaun Hooper, Linqing Wen, and Yanbei Chen. Towards low-latency real-time detection of gravitational waves from compact binary coalescences in the era of advanced detectors. *Phys. Rev. D*, 85:102002, 2012. doi: 10.1103/PhysRevD.85.102002.
- [238] Masahiro N. Machida. Binary Formation in Star-Forming Clouds with Various Metallicities. *Astrophys. J. Lett.*, 682:L1, 2008. doi: 10.1086/590109.
- [239] Piero Madau and Mark Dickinson. Cosmic star-formation history. *Annual Review of Astronomy and Astrophysics*, 52:415–486, 2014.
- [240] Piero Madau and Martin J. Rees. Massive black holes as Population III remnants. *Astrophys. J. Lett.*, 551:L27–L30, 2001. doi: 10.1086/319848.
- [241] Ignacio Magaña Hernandez. Measuring the polarization content of gravitational waves with strongly lensed binary black hole mergers. 11 2022.
- [242] Michele Maggiore, Chris Van Den Broeck, Nicola Bartolo, Enis Belgacem, Daniele Bertacca, Marie Anne Bizouard, Marica Branchesi, Sebastien Clesse, Stefano Foffa, Juan García-Bellido, et al. Science case for the einstein telescope. *Journal of Cosmology and Astroparticle Physics*, 2020(03):050, 2020.
- [243] K. G. Malmquist. On some relations in stellar statistics. *Meddelanden fran Lunds Astronomiska Observatorium Serie I*, 100:1–52, March 1922.
- [244] Pablo Marchant, Mathieu Renzo, Robert Farmer, Kaliroe MW Pappas, Ronald E Taam, Selma E De Mink, and Vassiliki Kalogera. Pulsational pair-instability supernovae in very close binaries. *The Astrophysical Journal*, 882(1):36, 2019.
- [245] Maxim Markevitch, A.H. Gonzalez, D. Clowe, A. Vikhlinin, L. David, W. Forman, C. Jones, S. Murray, and W. Tucker. Direct constraints on the dark matter self-interaction cross-section from the merging galaxy cluster 1E0657-56. *Astrophys. J.*, 606:819–824, 2004. doi: 10.1086/383178.
- [246] Maite Mateu-Lucena, Sascha Husa, Marta Colleoni, Héctor Estellés, Cecilio García-Quirós, David Keitel, Maria de Lluc Planas, and Antoni Ramos-Buades. Parameter estimation with the current generation of phenomenological waveform models applied to the black hole mergers of GWTC-1. *Mon. Not. Roy. Astron. Soc.*, 517(2):2403–2425, 2022. doi: 10.1093/mnras/stac2724.
- [247] Connor McIsaac, David Keitel, Thomas Collett, Ian Harry, Simone Mozzon, Oliver Edy, and David Bacon. Search for strongly lensed counterpart images of binary black hole mergers in the first two LIGO observing runs. *Phys. Rev. D*, 102(8):084031, 2020. doi: 10.1103/PhysRevD.102.084031.
- [248] Ashish Kumar Meena, Anuj Mishra, Anupreeta More, Sukanta Bose, and Jasjeet Singh Bagla. Gravitational lensing of gravitational waves: Probability

- of microlensing in galaxy-scale lens population. *Mon. Not. Roy. Astron. Soc.*, 517(1):872–884, 2022. doi: 10.1093/mnras/stac2721.
- [249] Cody Messick, Surabhi Sachdev, Kipp Cannon, Sarah Caudill, Chiwai Chan, Jolien D. E. Creighton, Ryan Everett, Becca Ewing, Heather Fong, Patrick Godwin, Chad Hanna, Rachael Huxford, Shasvath Kapadia, Alvin K. Y. Li, Rico K. L. Lo, Ryan Magee, Duncan Meacher, Siddharth R. Mohite, Debnandini Mukherjee, Atsushi Nishizawa, Hiroaki Ohta, Alexander Pace, Amit Reza, Minori Shikauchi, Leo Singer, Divya Singh, Javed Rana SK, Leo Tsukada, Daichi Tsuna, Takuya Tsutsui, Koh Ueno, and Aaron Zimmerman. Automating the Inclusion of Subthreshold Signal-to-Noise Ratios for Rapid Gravitational-Wave Localization. *arXiv e-prints*, art. arXiv:2011.02457, November 2020.
- [250] Cody Messick et al. Analysis Framework for the Prompt Discovery of Compact Binary Mergers in Gravitational-wave Data. *Phys. Rev. D*, 95(4):042001, 2017. doi: 10.1103/PhysRevD.95.042001.
- [251] R Benton Metcalf, MASSIMO Meneghetti, Camille Avestruz, Fabio Bel-lagamba, Clécio R Bom, Emmanuel Bertin, Rémi Cabanac, F Courbin, Andrew Davies, Etienne Decencière, et al. The strong gravitational lens finding challenge. *Astronomy & Astrophysics*, 625:A119, 2019.
- [252] John Michell. On the means of discovering the distance, magnitude, c. of the fixed stars, in consequence of the diminution of the velocity of their light, in case such a diminution should be found to take place in any of them, and such other data should be procured from observations, as would be farther necessary for that purpose. by the rev. john michell, b. d. f. r. s. in a letter to henry cavendish, esq. f. r. s. and a. s. *Philosophical Transactions of the Royal Society of London*, 74:35–57, 1784. ISSN 02610523. URL <http://www.jstor.org/stable/106576>.
- [253] Balazs Mikoczi, Matyas Vasuth, and Laszlo A. Gergely. Self-interaction spin effects in inspiralling compact binaries. *Phys. Rev. D*, 71:124043, 2005. doi: 10.1103/PhysRevD.71.124043.
- [254] M. Coleman Miller and Vanessa M. Lauburg. Mergers of Stellar-Mass Black Holes in Nuclear Star Clusters. *Astrophys. J.*, 692:917–923, 2009. doi: 10.1088/0004-637X/692/1/917.
- [255] Anuj Mishra. Introducing gwmat: Gravitational wave microlensing analysis tools. in prep.
- [256] Anuj Mishra, Ashish Kumar Meena, Anupreeta More, Sukanta Bose, and J S Bagla. Gravitational Lensing of Gravitational Waves: Effect of Microlens Population in Lensing Galaxies. *arXiv e-prints*, art. arXiv:2102.03946, February 2021.

- [257] Anuj Mishra, Ashish Kumar Meena, Anupreeta More, Sukanta Bose, and Jasjeet Singh Bagla. Gravitational lensing of gravitational waves: effect of microlens population in lensing galaxies. *Mon. Not. Roy. Astron. Soc.*, 508(4):4869–4886, 2021. doi: 10.1093/mnras/stab2875.
- [258] Chandra Kant Mishra, Aditya Kela, K. G. Arun, and Guillaume Faye. Ready-to-use post-Newtonian gravitational waveforms for binary black holes with nonprecessing spins: An update. *Phys. Rev. D*, 93(8):084054, 2016. doi: 10.1103/PhysRevD.93.084054.
- [259] Anupreeta More and Surhud More. Improved statistic to identify strongly lensed gravitational wave events. *Mon. Not. Roy. Astron. Soc.*, 515(1):1044–1051, 2022. doi: 10.1093/mnras/stac1704.
- [260] Anupreeta More, Remi Cabanac, Surhud More, Christophe Alard, Marceau Limousin, Jean-Paul Kneib, Raphael Gavazzi, and Veronica Motta. The cfhtls–strong lensing legacy survey (sl2s): Investigating the group-scale lenses with the sarcs sample. *The Astrophysical Journal*, 749(1):38, 2012.
- [261] Leonidas A. Moustakas, J. Brownstein, R. Fadely, C. D. Fassnacht, R. Gavazzi, T. Goodsall, R. L. Griffith, C. R. Keeton, J. P. Kneib, A. Koeke-moer, L. V. E. Koopmans, P. J. Marshall, J. Merten, R. B. Metcalf, M. Oguri, C. Papovich, H. Rein, R. Ryan, K. R. Stewart, and T. Treu. The Orphan Lenses Project. In *American Astronomical Society Meeting Abstracts #219*, volume 219 of *American Astronomical Society Meeting Abstracts*, page 146.01, January 2012.
- [262] Julian B. Muñoz, Ely D. Kovetz, Liang Dai, and Marc Kamionkowski. Lensing of Fast Radio Bursts as a Probe of Compact Dark Matter. *Phys. Rev. Lett.*, 117(9):091301, 2016. doi: 10.1103/PhysRevLett.117.091301.
- [263] Debnandini Mukherjee, Sarah Caudill, Ryan Magee, Cody Messick, Stephen Privitera, Surabhi Sachdev, Kent Blackburn, Patrick Brady, Patrick Brockill, Kipp Cannon, et al. Template bank for spinning compact binary mergers in the second observation run of advanced ligo and the first observation run of advanced virgo. *Physical Review D*, 103(8):084047, 2021.
- [264] Suvodip Mukherjee, Benjamin D. Wandelt, and Joseph Silk. Probing the theory of gravity with gravitational lensing of gravitational waves and galaxy surveys. *Mon. Not. Roy. Astron. Soc.*, 494(2):1956–1970, 2020. doi: 10.1093/mnras/staa827.
- [265] Suvodip Mukherjee, Benjamin D. Wandelt, and Joseph Silk. Multimessenger tests of gravity with weakly lensed gravitational waves. *Phys. Rev. D*, 101(10):103509, 2020. doi: 10.1103/PhysRevD.101.103509.
- [266] Suvodip Mukherjee, Tom Broadhurst, Jose M. Diego, Joseph Silk, and George F. Smoot. Inferring the lensing rate of LIGO-Virgo sources from

- the stochastic gravitational wave background. *Mon. Not. Roy. Astron. Soc.*, 501(2):2451–2466, 2021. doi: 10.1093/mnras/staa3813.
- [267] Suvodip Mukherjee, Tom Broadhurst, Jose M. Diego, Joseph Silk, and George F. Smoot. Impact of astrophysical binary coalescence time-scales on the rate of lensed gravitational wave events. *Mon. Not. Roy. Astron. Soc.*, 506(3):3751–3759, 2021. doi: 10.1093/mnras/stab1980.
- [268] Fumitaka Nakamura and Masayuki Umemura. On the Initial mass function of population III stars. *Astrophys. J.*, 548:19–32, 2001. doi: 10.1086/318663.
- [269] T. T. Nakamura and S. Deguchi. Wave Optics in Gravitational Lensing. *Progress of Theoretical Physics Supplement*, 133:137–153, January 1999. doi: 10.1143/PTPS.133.137.
- [270] Takahiro T. Nakamura. Gravitational lensing of gravitational waves from inspiraling binaries by a point mass lens. *Phys. Rev. Lett.*, 80:1138–1141, 1998. doi: 10.1103/PhysRevLett.80.1138.
- [271] Ramesh Narayan and Matthias Bartelmann. Lectures on gravitational lensing. In *13th Jerusalem Winter School in Theoretical Physics: Formation of Structure in the Universe Jerusalem, Israel, 27 December 1995 - 5 January 1996*, 1996.
- [272] Julio F Navarro, Carlos S Frenk, and Simon DM White. A universal density profile from hierarchical clustering. *The Astrophysical Journal*, 490(2):493, 1997.
- [273] Coenraad J. Neijssel, Alejandro Vigna-Gómez, Simon Stevenson, Jim W. Barrett, Sebastian M. Gaebel, Floor Broekgaarden, Selma E. de Mink, Dorottya Szécsi, Serena Vinciguerra, and Ilya Mandel. The effect of the metallicity-specific star formation history on double compact object mergers. *Mon. Not. Roy. Astron. Soc.*, 490(3):3740–3759, 2019. doi: 10.1093/mnras/stz2840.
- [274] Ken K.Y. Ng, Kaze W.K. Wong, Tom Broadhurst, and Tjonnie G.F. Li. Precise LIGO Lensing Rate Predictions for Binary Black Holes. *Phys. Rev. D*, 97(2):023012, 2018. doi: 10.1103/PhysRevD.97.023012.
- [275] Philippe Nguyen et al. Environmental noise in advanced LIGO detectors. *Class. Quant. Grav.*, 38(14):145001, 2021. doi: 10.1088/1361-6382/ac011a.
- [276] Alexander H. Nitz, Thomas Dent, Tito Dal Canton, Stephen Fairhurst, and Duncan A. Brown. Detecting binary compact-object mergers with gravitational waves: Understanding and Improving the sensitivity of the PyCBC search. *Astrophys. J.*, 849(2):118, 2017. doi: 10.3847/1538-4357/aa8f50.
- [277] Alexander H. Nitz, Tito Dal Canton, Derek Davis, and Steven Reyes. Rapid detection of gravitational waves from compact binary mergers with PyCBC Live. *Phys. Rev. D*, 98(2):024050, 2018. doi: 10.1103/PhysRevD.98.024050.

- [278] Alexander H. Nitz, Collin D. Capano, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Marlin Schäfer, Rahul Dhurkunde, and Miriam Cabero. 3-OGC: Catalog of Gravitational Waves from Compact-binary Mergers. *Astrophys. J.*, 922(1):76, 2021. doi: 10.3847/1538-4357/ac1c03.
- [279] Alexander H. Nitz et al. gwastro/pycbc: PyCBC Release v1.15.3, 2019.
- [280] Michael L. Norman. Population III Star Formation and IMF. *AIP Conf. Proc.*, 990(1):3–15, 2008. doi: 10.1063/1.2905629.
- [281] Masamune Oguri. Effect of gravitational lensing on the distribution of gravitational waves from distant binary black hole mergers. *Mon. Not. Roy. Astron. Soc.*, 480(3):3842–3855, 2018. doi: 10.1093/mnras/sty2145.
- [282] Masamune Oguri and Philip J. Marshall. Gravitationally lensed quasars and supernovae in future wide-field optical imaging surveys. *Mon. Not. Roy. Astron. Soc.*, 405:2579–2593, 2010. doi: 10.1111/j.1365-2966.2010.16639.x.
- [283] Masamune Oguri and Ryuichi Takahashi. Probing Dark Low-mass Halos and Primordial Black Holes with Frequency-dependent Gravitational Lensing Dispersions of Gravitational Waves. *Astrophys. J.*, 901(1):58, 2020. doi: 10.3847/1538-4357/abafab.
- [284] H.C. Ohanian. On the focusing of gravitational radiation. *Int. J. Theor. Phys.*, 9:425–437, 1974. doi: 10.1007/BF01810927.
- [285] Serguei Ossokine et al. Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation. *Phys. Rev. D*, 102(4):044055, 2020. doi: 10.1103/PhysRevD.102.044055.
- [286] Giulia Pagano, Otto A. Hannuksela, and Tjonnie G. F. Li. lensingGW: a Python package for lensing of gravitational waves. *Astron. Astrophys.*, 643:A167, 2020. doi: 10.1051/0004-6361/202038730.
- [287] Carlos Palenzuela. Introduction to Numerical Relativity. *Front. Astron. Space Sci.*, 7:58, 2020. doi: 10.3389/fspas.2020.00058.
- [288] Peter TH Pang, Otto A Hannuksela, Tim Dietrich, Giulia Pagano, and Ian W Harry. Lensed or not lensed: Determining lensing magnifications for binary neutron star mergers from a single detection. *Monthly Notices of the Royal Astronomical Society*, 495(4):3740–3750, 2020.
- [289] C. Pankow, P. Brady, E. Ochsner, and R. O’Shaughnessy. Novel scheme for rapid parallel parameter estimation of gravitational waves from compact binary coalescences. *Phys. Rev. D*, 92(2):023002, 2015. doi: 10.1103/PhysRevD.92.023002.

- [290] Danuta Paraficz, Jens Hjorth, and Ardis Eliasdottir. Results of optical monitoring of 5 SDSS double QSOs with the Nordic Optical Telescope. *Astron. Astrophys.*, 499:395, 2009. doi: 10.1051/0004-6361/200811387.
- [291] Ethan Payne, Sophie Hourihane, Jacob Golomb, Rhiannon Udall, Richard Udall, Derek Davis, and Katerina Chatziioannou. Curious case of GW200129: Interplay between spin-precession inference and data-quality issues. *Phys. Rev. D*, 106(10):104017, 2022. doi: 10.1103/PhysRevD.106.104017.
- [292] James Paynter, Rachel Webster, and Eric Thrane. Evidence for an intermediate-mass black hole from a gravitationally lensed gamma-ray burst. *Nature Astron.*, 5(6):560–568, 2021. doi: 10.1038/s41550-021-01307-1.
- [293] Eric Poisson. Gravitational waves from inspiraling compact binaries: The Quadrupole moment term. *Phys. Rev. D*, 57:5287–5290, 1998. doi: 10.1103/PhysRevD.57.5287.
- [294] Geraint Pratten, Sascha Husa, Cecilio Garcia-Quiros, Marta Colleoni, Antoni Ramos-Buades, Hector Estelles, and Rafel Jaume. Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes. *Phys. Rev. D*, 102(6):064001, 2020. doi: 10.1103/PhysRevD.102.064001.
- [295] Geraint Pratten et al. Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes. *Phys. Rev. D*, 103(10):104056, 2021. doi: 10.1103/PhysRevD.103.104056.
- [296] M. Punturo et al. The third generation of gravitational wave observatories and their science reach. *Class. Quant. Grav.*, 27:084007, 2010. doi: 10.1088/0264-9381/27/8/084007.
- [297] Michael Pürrer. Frequency domain reduced order models for gravitational waves from aligned-spin compact binaries. *Class. Quant. Grav.*, 31(19):195010, 2014. doi: 10.1088/0264-9381/31/19/195010.
- [298] Michael Pürrer. Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass-ratios and spins. *Phys. Rev. D*, 93(6):064041, 2016. doi: 10.1103/PhysRevD.93.064041.
- [299] S. Rahvar and J. W. Moffat. Propagation of Electromagnetic Waves in MOG: Gravitational Lensing. *Mon. Not. Roy. Astron. Soc.*, 482:4514, 2019. doi: 10.1093/mnras/sty3002.
- [300] David Reitze et al. Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO. *Bull. Am. Astron. Soc.*, 51(7):035, 2019.

- [301] E. Ritondale, M. W. Auger, S. Vegetti, and J. P. McKean. Resolving on 100 pc scales the UV-continuum in Lyman- α emitters between redshift 2 and 3 with gravitational lensing. *Monthly Notices of the Royal Astronomical Society*, 482(4):4744–4762, February 2019. doi: 10.1093/mnras/sty2833.
- [302] Andrew Robertson, Graham P. Smith, Richard Massey, Vincent Eke, Mathilde Jauzac, Matteo Bianconi, and Dan Ryczanowski. What does strong gravitational lensing? The mass and redshift distribution of high-magnification lenses. *Mon. Not. Roy. Astron. Soc.*, 495(4):3727–3739, July 2020. doi: 10.1093/mnras/staa1429.
- [303] Florent Robinet, Nicolas Arnaud, Nicolas Leroy, Andrew Lundgren, Duncan Macleod, and Jessica McIver. Omicron: a tool to characterize transient noise in gravitational-wave detectors. *SoftwareX*, 12:100620, 2020. doi: 10.1016/j.softx.2020.100620.
- [304] Carl L. Rodriguez and Abraham Loeb. Redshift Evolution of the Black Hole Merger Rate from Globular Clusters. *Astrophys. J. Lett.*, 866(1):L5, 2018. doi: 10.3847/2041-8213/aae377.
- [305] I. M. Romero-Shaw et al. Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue. *Mon. Not. Roy. Astron. Soc.*, 499(3):3295–3319, 2020. doi: 10.1093/mnras/staa2850.
- [306] Javier Roulet, Tejaswi Venumadhav, Barak Zackay, Liang Dai, and Matias Zaldarriaga. Binary Black Hole Mergers from LIGO/Virgo O1 and O2: Population Inference Combining Confident and Marginal Events. *Phys. Rev. D*, 102(12):123022, 2020. doi: 10.1103/PhysRevD.102.123022.
- [307] Vera C. Rubin and Jr. Ford, W. Kent. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *The Astrophysical Journal*, 159: 379, February 1970. doi: 10.1086/150317.
- [308] Dan Ryczanowski, Graham P. Smith, Matteo Bianconi, Richard Massey, Andrew Robertson, and Mathilde Jauzac. On building a cluster watch-list for identifying strongly lensed supernovae, gravitational waves and kilonovae. *Mon. Not. Roy. Astron. Soc.*, 495(2):1666–1671, 2020. doi: 10.1093/mnras/staa1274.
- [309] Surabhi Sachdev et al. The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO’s Second and Advanced Virgo’s First Observing Runs. 1 2019.
- [310] Shio Sakon et al. Template bank for compact binary mergers in the fourth observing run of Advanced LIGO, Advanced Virgo, and KAGRA. *Phys. Rev. D*, 109(4):044066, 2024. doi: 10.1103/PhysRevD.109.044066.

- [311] Filippo Santoliquido, Michela Mapelli, Nicola Giacobbo, Yann Bouffanais, and M. Celeste Artale. The cosmic merger rate density of compact objects: impact of star formation, metallicity, initial mass function and binary evolution. *Mon. Not. Roy. Astron. Soc.*, 502(4):4877–4889, 2021. doi: 10.1093/mnras/stab280.
- [312] B. S. Sathyaprakash and S. V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 44:3819–3834, 1991. doi: 10.1103/PhysRevD.44.3819.
- [313] Daniel Schaerer. On the Properties of massive population III stars and metal-free stellar populations. *Astron. Astrophys.*, 382:28–42, 2002. doi: 10.1051/0004-6361:20011619.
- [314] Patricia Schmidt. Gravitational Waves From Binary Black Hole Mergers: Modeling and Observations. *Front. Astron. Space Sci.*, 7:28, 2020. doi: 10.3389/fspas.2020.00028.
- [315] Peter Schneider, Jürgen Ehlers, and Emilio E. Falco. *Gravitational Lenses*. 1992. doi: 10.1007/978-3-662-03758-4.
- [316] Eungwang Seo, Otto A. Hannuksela, and Tjonnie G. F. Li. Improving Detection of Gravitational-wave Microlensing Using Repeated Signals Induced by Strong Lensing. *Astrophys. J.*, 932(1):50, 2022. doi: 10.3847/1538-4357/ac6dea.
- [317] M. Sereno, Ph. Jetzer, A. Sesana, and M. Volonteri. Cosmography with strong lensing of LISA gravitational wave sources. *Mon. Not. Roy. Astron. Soc.*, 415:2773, 2011. doi: 10.1111/j.1365-2966.2011.18895.x.
- [318] Yiping Shu, Adam S. Bolton, Shude Mao, Christopher S. Kochanek, Ismael Pérez-Fournon, Masamune Oguri, Antonio D. Montero-Dorta, Matthew A. Cornachione, Rui Marques-Chaves, Zheng Zheng, Joel R. Brownstein, and Brice Ménard. The BOSS Emission-line Lens Survey. IV. Smooth Lens Models for the BELLS GALLERY Sample. *The Astrophysical Journal*, 833(2):264, December 2016. doi: 10.3847/1538-4357/833/2/264.
- [319] Leo P. Singer and Larry R. Price. Rapid Bayesian position reconstruction for gravitational-wave transients. *Phys. Rev. D*, 93(2):024013, 2016. doi: 10.1103/PhysRevD.93.024013.
- [320] Leo P. Singer, Daniel A. Goldstein, and Joshua S. Bloom. The Two LIGO/Virgo Binary Black Hole Mergers on 2019 August 28 Were Not Strongly Lensed. *arXiv e-prints*, art. arXiv:1910.03601, October 2019. doi: 10.48550/arXiv.1910.03601.
- [321] John Skilling. Nested sampling for general Bayesian computation. *Bayesian Analysis*, 1(4):833–859, 2006. doi: 10.1214/06-BA127.

- [322] G. P. Smith, M. Bianconi, R. Massey, and A. Robertson. LIGO/Virgo S190426c: a possible lensed NS-NS merger. *GRB Coordinates Network*, 24250:1, January 2019.
- [323] G.P. Smith et al. Strong-lensing of Gravitational Waves by Galaxy Clusters. *IAU Symp.*, 338:98–102, 2017. doi: 10.1017/S1743921318003757.
- [324] Graham P. Smith, Mathilde Jauzac, John Veitch, Will M. Farr, Richard Massey, and Johan Richard. What if LIGO’s gravitational wave detections are strongly lensed by massive galaxy clusters? *Mon. Not. Roy. Astron. Soc.*, 475(3):3823–3828, 2018. doi: 10.1093/mnras/sty031.
- [325] Graham P. Smith, Andrew Robertson, Matteo Bianconi, and Mathilde Jauzac. Discovery of Strongly-lensed Gravitational Waves - Implications for the LSST Observing Strategy. *arXiv e-prints*, art. arXiv:1902.05140, February 2019. doi: 10.48550/arXiv.1902.05140.
- [326] Rory J. E. Smith, Gregory Ashton, Avi Vajpeyi, and Colm Talbot. Massively parallel Bayesian inference for transient gravitational-wave astronomy. *Mon. Not. Roy. Astron. Soc.*, 498(3):4492–4502, 2020. doi: 10.1093/mnras/staa2483.
- [327] S. Soni et al. Reducing scattered light in LIGO’s third observing run. *Class. Quant. Grav.*, 38(2):025016, 2020. doi: 10.1088/1361-6382/abc906.
- [328] S. Soni et al. Discovering features in gravitational-wave data through detector characterization, citizen science and machine learning. *Class. Quant. Grav.*, 38(19):195016, 2021. doi: 10.1088/1361-6382/ac1ccb.
- [329] Alessandro Sonnenfeld, Raphaël Gavazzi, Sherry H. Suyu, Tommaso Treu, and Philip J. Marshall. The SL2S Galaxy-scale Lens Sample. III. Lens Models, Surface Photometry, and Stellar Masses for the Final Sample. *The Astrophysical Journal*, 777(2):97, November 2013. doi: 10.1088/0004-637X/777/2/97.
- [330] Alessandro Sonnenfeld, James H H Chan, Yiping Shu, Anupreeta More, Masamune Oguri, Sherry H Suyu, Kenneth C Wong, Chien-Hsiu Lee, Jean Coupon, Atsunori Yonehara, Adam S Bolton, Anton T Jaelani, Masayuki Tanaka, Satoshi Miyazaki, and Yutaka Komiyama. Survey of Gravitationally-lensed Objects in HSC Imaging (SuGOHI). I. Automatic search for galaxy-scale strong lenses. *Publications of the Astronomical Society of Japan*, 70 (SP1), 09 2017. ISSN 0004-6264. doi: 10.1093/pasj/psx062. URL <https://doi.org/10.1093/pasj/psx062>. S29.
- [331] Sonnenfeld, Alessandro, Jaelani, Anton T., Chan, James, More, Anupreeta, Suyu, Sherry H., Wong, Kenneth C., Oguri, Masamune, and Lee, Chien-Hsiu. Survey of gravitationally-lensed objects in hsc imaging (sugohi) - iii. statistical strong lensing constraints on the stellar imf of cmass galaxies. *Astronomy*

- & *Astrophysics*, 630:A71, 2019. doi: 10.1051/0004-6361/201935743. URL <https://doi.org/10.1051/0004-6361/201935743>.
- [332] Sonnenfeld, Alessandro, Verma, Aprajita, More, Anupreeta, Baeten, Elisabeth, Macmillan, Christine, Wong, Kenneth C., Chan, James H. H., Jaelani, Anton T., Lee, Chien-Hsiu, Oguri, Masamune, Rusu, Cristian E., Veldthuis, Marten, Trouille, Laura, Marshall, Philip J., Hutchings, Roger, Allen, Campbell, O' Donnell, James, Cornen, Claude, Davis, Christopher P., McMaster, Adam, Lintott, Chris, and Miller, Grant. Survey of gravitationally-lensed objects in hsc imaging (sugohi) - vi. crowdsourced lens finding with space warps. *Astronomy & Astrophysics*, 642:A148, 2020. doi: 10.1051/0004-6361/202038067. URL <https://doi.org/10.1051/0004-6361/202038067>.
- [333] Joshua S Speagle. dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences. *Monthly Notices of the Royal Astronomical Society*, 493(3):3132–3158, 02 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa278. URL <https://doi.org/10.1093/mnras/staa278>.
- [334] Justin Spilker, Daniel Marrone, Manuel Aravena, M. Bethermin, M. Bothwell, John Carlstrom, Scott Chapman, Tom Crawford, Carlos Breuck, C. Fassnacht, Anthony Gonzalez, Thomas Greve, Yashar Hezaveh, Katrina Litke, Jingzhe Ma, Manny Malkan, Kaja Rotermund, Maria Strandet, Joaquin Vieira, and Niraj Welikala. Alma imaging and gravitational lens models of south pole telescope-selected dusty, star-forming galaxies at high redshifts. *The Astrophysical Journal*, 826, 04 2016. doi: 10.3847/0004-637X/826/2/112.
- [335] Ling Sun, Evan Goetz, Jeffrey S. Kissel, Joseph Betzwieser, Sudarshan Karki, Dripta Bhattacharjee, Pep B. Covas, Laurence E. H. Datrier, Shivaraj Kandhasamy, Yannick K. Lecoecue, Gregory Mendell, Timesh Mistry, Ethan Payne, Richard L. Savage, Aaron Viets, Madeline Wade, Alan J. Weinstein, Stuart Aston, Craig Cahillane, Jennifer C. Driggers, Sheila E. Dwyer, and Alexander Urban. Characterization of systematic error in Advanced LIGO calibration in the second half of O3. *arXiv e-prints*, art. arXiv:2107.00129, June 2021. doi: 10.48550/arXiv.2107.00129.
- [336] Ling Sun et al. Characterization of systematic error in Advanced LIGO calibration. *Class. Quant. Grav.*, 37(22):225008, 2020. doi: 10.1088/1361-6382/abb14e.
- [337] Ryuichi Takahashi and Takashi Nakamura. Wave effects in gravitational lensing of gravitational waves from chirping binaries. *Astrophys. J.*, 595: 1039–1051, 2003. doi: 10.1086/377430.
- [338] Colm Talbot, Rory Smith, Eric Thrane, and Gregory B. Poole. Parallelized Inference for Gravitational-Wave Astronomy. *Phys. Rev. D*, 100(4):043030, 2019. doi: 10.1103/PhysRevD.100.043030.

- [339] Giovanni Tambalo, Miguel Zumalacárregui, Liang Dai, and Mark Ho-Yeuk Cheung. Gravitational wave lensing as a probe of halo properties and dark matter. *Phys. Rev. D*, 108(10):103529, 2023. doi: 10.1103/PhysRevD.108.103529.
- [340] K.S. Thorne. THE THEORY OF GRAVITATIONAL RADIATION: AN INTRODUCTORY REVIEW. In *Les Houches Summer School on Gravitational Radiation*, pages 1–57, 1 1982.
- [341] Eric Thrane and Colm Talbot. An introduction to Bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models. *Publications of the Astronomical Society of Australia*, 36:e010, March 2019. doi: 10.1017/pasa.2019.2.
- [342] Jeremy L. Tinker, Andrey V. Kravtsov, Anatoly Klypin, Kevork Abazajian, Michael S. Warren, Gustavo Yepes, Stefan Gottlober, and Daniel E. Holz. Toward a halo mass function for precision cosmology: The Limits of universality. *Astrophys. J.*, 688:709–728, 2008. doi: 10.1086/591439.
- [343] Arthur E. Tolley, Gareth S. Cabourn Davies, Ian W. Harry, and Andrew P. Lundgren. ArchEnemy: removing scattered-light glitches from gravitational wave data. *Class. Quant. Grav.*, 40(16):165005, 2023. doi: 10.1088/1361-6382/ace22f.
- [344] Paul Torrey, Sarah Wellons, Francisco Machado, Brendan Griffen, Dylan Nelson, Vicente Rodriguez-Gomez, Ryan McKinnon, Annalisa Pillepich, Chung-Pei Ma, Mark Vogelsberger, Volker Springel, and Lars Hernquist. An analysis of the evolving comoving number density of galaxies in hydrodynamical simulations. *Monthly Notices of the Royal Astronomical Society*, 454 (3):2770–2786, 10 2015. ISSN 0035-8711. doi: 10.1093/mnras/stv1986. URL <https://doi.org/10.1093/mnras/stv1986>.
- [345] Kim-Vy H. Tran, Anishya Harshan, Karl Glazebrook, G. C. Keerthi Vasan, Tucker Jones, Colin Jacobs, Glenn G. Kacprzak, Tania M. Barone, Thomas E. Collett, Anshu Gupta, Astrid Henderson, Lisa J. Kewley, Sebastian Lopez, Themiya Nanayakkara, Ryan L. Sanders, and Sarah M. Sweet. The AGEL Survey: Spectroscopic Confirmation of Strong Gravitational Lenses in the DES and DECaLS Fields Selected Using Convolutional Neural Networks. *The Astronomical Journal*, 164(4):148, October 2022. doi: 10.3847/1538-3881/ac7da2.
- [346] Tommaso Treu, Leon Koopmans, Adam Bolton, Scott Burles, and Leonidas Moustakas. The sloan-lens acs survey. 2. stellar populations and internal structure of early-type lens galaxies. *Astrophys. J.*, 640:662–672, 2006. doi: 10.1086/500124.

- [347] Takuya Tsutsui, Kipp Cannon, and Leo Tsukada. High speed source localization in searches for gravitational waves from compact object collisions. *Phys. Rev. D*, 103(4):043011, 2021. doi: 10.1103/PhysRevD.103.043011.
- [348] Slava G. Turyshev and Viktor T. Toth. Diffraction of electromagnetic waves in the gravitational field of the Sun. *Phys. Rev.*, D96(2):024008, 2017. doi: 10.1103/PhysRevD.96.024008.
- [349] Richard Udall, Joshua Brandt, Grihith Manchanda, Adhav Arulanandan, James Clark, Jacob Lange, Richard O’Shaughnessy, and Laura Cadonati. RUNMON-RIFT: Adaptive configuration and healing for large-scale parameter inference. *Astron. Comput.*, 42:100664, 2023. doi: 10.1016/j.ascom.2022.100664.
- [350] Juan Urrutia and Ville Vaskonen. Lensing of gravitational waves as a probe of compact dark matter. *Mon. Not. Roy. Astron. Soc.*, 509(1):1358–1365, 2021. doi: 10.1093/mnras/stab3118.
- [351] Samantha A. Usman et al. The PyCBC search for gravitational waves from compact binary coalescence. *Class. Quant. Grav.*, 33(21):215004, 2016. doi: 10.1088/0264-9381/33/21/215004.
- [352] Gabriele Vajente, Yiwen Huang, Maximiliano Isi, Jenne C. Driggers, Jeffrey S. Kissel, Marek J. Szczepanczyk, and Salvatore Vitale. Machine-learning nonstationary noise out of gravitational-wave detectors. *Phys. Rev. D*, 101(4):042003, 2020. doi: 10.1103/PhysRevD.101.042003.
- [353] Michele Vallisneri. Testing general relativity with gravitational waves: a reality check. *Phys. Rev. D*, 86:082001, 2012. doi: 10.1103/PhysRevD.86.082001.
- [354] Vijay Varma, Scott E. Field, Mark A. Scheel, Jonathan Blackman, Davide Gerosa, Leo C. Stein, Lawrence E. Kidder, and Harald P. Pfeiffer. Surrogate models for precessing binary black hole simulations with unequal masses. *Phys. Rev. Research.*, 1:033015, 2019. doi: 10.1103/PhysRevResearch.1.033015.
- [355] H. K. Vedantham et al. Symmetric Achromatic Variability in Active Galaxies – A Powerful New Gravitational Lensing Probe? *Astrophys. J.*, 845(2):89, 2017. doi: 10.3847/1538-4357/aa745c.
- [356] J. Veitch et al. Robust parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys. Rev. D*, 91:042003, Feb 2015. doi: 10.1103/PhysRevD.91.042003.
- [357] Aaron Viets et al. Reconstructing the calibrated strain signal in the Advanced LIGO detectors. *Class. Quant. Grav.*, 35(9):095015, 2018. doi: 10.1088/1361-6382/aab658.

- [358] Aditya Vijaykumar, Ajit Kumar Mehta, and Apratim Ganguly. Detection and parameter estimation challenges of type-II lensed binary black hole signals. *Phys. Rev. D*, 108(4):043036, 2023. doi: 10.1103/PhysRevD.108.043036.
- [359] Digvijay Wadekar, Tejaswi Venumadhav, Ajit Kumar Mehta, Javier Roulet, Seth Olsen, Jonathan Mushkin, Barak Zackay, and Matias Zaldarriaga. A new approach to template banks of gravitational waves with higher harmonics: reducing matched-filtering cost by over an order of magnitude. *arXiv e-prints*, art. arXiv:2310.15233, October 2023. doi: 10.48550/arXiv.2310.15233.
- [360] Yijun Wang, Rico K. L. Lo, Alvin K. Y. Li, and Yanbei Chen. Identifying Type II Strongly Lensed Gravitational-Wave Images in Third-Generation Gravitational-Wave Detectors. *Phys. Rev. D*, 103(10):104055, 2021. doi: 10.1103/PhysRevD.103.104055.
- [361] Yun Wang, Albert Stebbins, and Edwin L. Turner. Gravitational lensing of gravitational waves from merging neutron star binaries. *Phys. Rev. Lett.*, 77: 2875–2878, 1996. doi: 10.1103/PhysRevLett.77.2875.
- [362] A. Weiss et al. ALMA redshifts of millimeter-selected galaxies from the SPT survey: The redshift distribution of dusty star-forming galaxies. *Astrophys. J.*, 767:88, 2013. doi: 10.1088/0004-637X/767/1/88.
- [363] Ewoud Wempe, Léon V. E. Koopmans, A. Renske A. C. Wierda, Otto Akseli Hannuksela, and Chris van den Broeck. A lensing multi-messenger channel: Combining ligo-virgo-kagra lensed gravitational-wave measurements with euclid observations. *arXiv preprint arXiv:2204.08732*, 2022.
- [364] A. Renske A. C. Wierda, Ewoud Wempe, Otto A. Hannuksela, Léon V. E. Koopmans, and Chris Van Den Broeck. Beyond the Detector Horizon: Forecasting Gravitational-Wave Strong Lensing. *Astrophys. J.*, 921(2):154, 2021. doi: 10.3847/1538-4357/ac1bb4.
- [365] H. J. Witt. Investigation of high amplification events in light curves of gravitationally lensed quasars. *Astronomy & Astrophysics*, 236:311, September 1990.
- [366] Henry W. Y. Wong, Lok W. L. Chan, Isaac C. F. Wong, Rico K. L. Lo, and Tjonnie G. F. Li. Using overlap of sky localization probability maps for filtering potentially lensed pairs of gravitational-wave signals. *arXiv e-prints*, art. arXiv:2112.05932, December 2021. doi: 10.48550/arXiv.2112.05932.
- [367] Kaze W. K. Wong, Katelyn Breivik, Kyle Kremer, and Thomas Callister. Joint constraints on the field-cluster mixing fraction, common envelope efficiency, and globular cluster radii from a population of binary hole mergers via deep learning. *Phys. Rev. D*, 103(8):083021, 2021. doi: 10.1103/PhysRevD.103.083021.

- [368] Kenneth C. Wong, Alessandro Sonnenfeld, James H. H. Chan, Cristian E. Rusu, Masayuki Tanaka, Anton T. Jaelani, Chien-Hsiu Lee, Anupreeta More, Masamune Oguri, Sherry H. Suyu, and Yutaka Komiyama. Survey of gravitationally lensed objects in hsc imaging (sugohi). ii. environments and line-of-sight structure of strong gravitational lens galaxies to $z \sim 0.8$. *The Astrophysical Journal*, 867(2):107, nov 2018. doi: 10.3847/1538-4357/aae381. URL <https://dx.doi.org/10.3847/1538-4357/aae381>.
- [369] Kenneth C Wong, James H H Chan, Dani C-Y Chao, Anton T Jaelani, Issha Kayo, Chien-Hsiu Lee, Anupreeta More, and Masamune Oguri. Survey of Gravitationally lensed objects in HSC Imaging (SuGOHI). VIII. New galaxy-scale lenses from the HSC SSP. *Publications of the Astronomical Society of Japan*, 74(5):1209–1219, 08 2022. ISSN 0004-6264. doi: 10.1093/pasj/psac065. URL <https://doi.org/10.1093/pasj/psac065>.
- [370] Mick Wright and Martin Hendry. Gravelamps: Gravitational Wave Lensing Mass Profile Model Selection. *The Astrophysical Journal*, 935(2):68, August 2022. doi: 10.3847/1538-4357/ac7ec2.
- [371] D. Wysocki, R. O’Shaughnessy, Jacob Lange, and Yao-Lung L. Fang. Accelerating parameter inference with graphics processing units. *Phys. Rev. D*, 99(8):084026, 2019. doi: 10.1103/PhysRevD.99.084026.
- [372] Fei Xu, Jose Maria Ezquiaga, and Daniel E. Holz. Please Repeat: Strong Lensing of Gravitational Waves as a Probe of Compact Binary and Galaxy Populations. *Astrophys. J.*, 929(1):9, 2022. doi: 10.3847/1538-4357/ac58f8.
- [373] Simon M. C. Yeung, Mark H. Y. Cheung, Eungwang Seo, Joseph A. J. Gais, Otto A. Hannuksela, and Tjonnie G. F. Li. Detectability of microlensed gravitational waves. *Mon. Not. Roy. Astron. Soc.*, 526(2):2230–2240, 2023. doi: 10.1093/mnras/stad2772.
- [374] Hai Yu, Pengjie Zhang, and Fa-Yin Wang. Strong lensing as a giant telescope to localize the host galaxy of gravitational wave event. *Mon. Not. Roy. Astron. Soc.*, 497(1):204–209, 2020. doi: 10.1093/mnras/staa1952.
- [375] Michael Zevin, Simone S. Bavera, Christopher P. L. Berry, Vicky Kalogera, Tassos Fragos, Pablo Marchant, Carl L. Rodriguez, Fabio Antonini, Daniel E. Holz, and Chris Pankow. One Channel to Rule Them All? Constraining the Origins of Binary Black Holes Using Multiple Formation Pathways. *Astrophys. J.*, 910(2):152, 2021. doi: 10.3847/1538-4357/abe40e.
- [376] Huan Zhou, Zhengxiang Li, Kai Liao, and Zhiqi Huang. Constraints on compact dark matter from lensing of gravitational waves for the third-generation gravitational wave detector. *Mon. Not. Roy. Astron. Soc.*, 518(1):149–156, 2022. doi: 10.1093/mnras/stac2944.

- [377] Miguel Zumalacarregui and Uros Seljak. Limits on stellar-mass compact objects as dark matter from gravitational lensing of type Ia supernovae. *Phys. Rev. Lett.*, 121(14):141101, 2018. doi: 10.1103/PhysRevLett.121.141101.
- [378] F. Zwicky. Nebulae as gravitational lenses. *Phys. Rev.*, 51:290–290, Feb 1937. doi: 10.1103/PhysRev.51.290. URL <https://link.aps.org/doi/10.1103/PhysRev.51.290>.

Before I end my thesis, I would like to include the last part of the song “A Winter Story”.

あなたの全てが
形を無くしても
永遠に僕の中で生きてくよ
さようなら出来ずに
歩き出す僕と
ずっと一緒に

Even though your everything has lost
its shape, you will still live inside me
eternally. Although I was unable to say
goodbye, I will now move on and con-
tinue with my life, for I know you will
be with me forever.

Dear mother, this thesis and my thesis defense mark the end of my PhD journey and the end of a chapter in my life. While I will never be able to speak to you again, you will continue to live inside my heart forever. I promise I will never forget about you, but for now, I am going to move on to the next stage of my life. May you rest in peace.