

Bayesian Inference for Fast Scattering Glitches

Aislinn D. McCann
California State University, Northridge

Mentors: Rhiannon Udall and Derek Davis
LIGO, California Institute of Technology, Pasadena

Final Report
LIGO Caltech SURF Program 2023

Data collected by gravitational wave (GW) interferometers such as the Laser Interferometer Gravitational-wave Observatory (LIGO) is permeated by noise as a result of environmental interference. Parameter estimation pipelines such as `Bilby` used to analyse LIGO data employs Bayesian inference, which assumes that the noise in GW data is Gaussian and stationary: an assumption contradicted by the nature of non-Gaussian transient noise “glitches” prevalent within the data. We have constructed a mathematical model that emulates the waveform of fast scattering glitches, which we tested via `Bilby` to determine the efficacy of glitch mitigation under the basis of the model. The implementation of this model will facilitate the efficient subtraction of real fast scattering glitch instances from GW strain data, allowing for improved analysis and signal detection for future observing runs.

I. INTRODUCTION

The Laser Interferometer Gravitational-wave Observatory (LIGO) is an observatory designed to detect gravitational waves (GWs) via the use of two separate Michelson interferometer detectors. Within each detector is a beam splitter which sends a laser through each arm, allowing the split beam on either side to cycle in a Fabry-Perot cavity and rejoin the main beam to be analysed by a photodetector [1]. When GWs pass through a detector, each arm experiences a slight displacement which creates instances of constructive or destructive interference from the recombined beam, thereby inducing a phase shift which is then converted into a measurable signal [2]. A high sensitivity is required for GW detectors to receive data from distant sources such as compact binary coalescences (CBCs). The characteristic LIGO strain, lying at $10^{-21} \text{ Hz}^{-1/2}$, is easily overpowered by persistent and short duration transient noise “glitches” produced by various sources of environmental interference or electronic malfunction which rise above the Gaussian noise floor [2–5].

One form of glitch, known as scattered light glitches, arises when segments of the laser diverge from the main beam path. These segments reflect off of objects with non-zero relative velocities within the interferometer, which later rejoin the main beam and produce an additional phase shift [4, 6]. Scattered light glitches introduce two primary challenges in analysis. First, glitch instances may trigger false positives in GW search pipelines. The second and more frequent complication is that glitches may overlap on top of an existing signal, posing the largest hindrance to analysis efforts. Our focus is to mitigate instances of fast scattering glitches, a form of scattering glitch which occurs as a result of increased ground activity in the anthropogenic band (1 – 5 Hz) and microseism band (0.1 – 0.3 Hz). The scattering mechanism

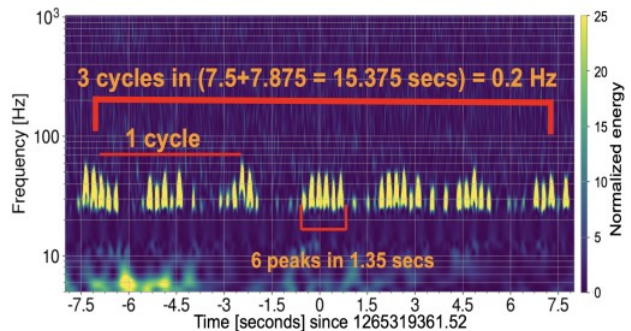


FIG. 1: A spectrogram of fast scattering triggers generated using the Q-transform. Fast scattering glitches are characterised by multiple peaks organised in a formation resembling a broad arch. Image reproduced from [2].

attributed to these sources couples into the differential arm motion, or DARM, and affects the detector’s sensitivity in the frequency band between 10 and 50 Hz [2]. Figure 1 provides an example of the short duration noise bursts characteristic of fast scattering glitches.

Removing noise and glitches from GW strain data has been a persistent effort by the GW scientific community to improve detector sensitivity. It is not only necessary to produce a reliable result, but it is also an imperative step to improve the accuracy of CBC parameter estimation pipelines which analyse raw strain data collected by detectors to infer astrophysical properties that characterise GW sources [4]. One such pipeline is `Bilby`, a Python library which utilises Bayesian inference to perform accurate parameter estimations [7].

Bayesian inference incorporates Bayes’ theorem to produce the posterior probability distribution of GW source parameters by incorporating the prior distribution of these source parameters with a model hypothesis. The

posterior probability may be computed using Bayes' theorem with data d and source parameters θ given a model hypothesis \mathcal{M} [3, 8]:

$$p(\theta|d, \mathcal{M}) = \frac{\mathcal{L}(d|\theta, \mathcal{M})\pi(\theta|\mathcal{M})}{\mathcal{Z}(d|\mathcal{M})}, \quad (1)$$

where $\mathcal{L}(d|\theta, \mathcal{M})$ is the likelihood and $\pi(\theta|\mathcal{M})$ is the prior probability. The evidence $\mathcal{Z}(d|\mathcal{M})$, a normalisation constant, is described by

$$\mathcal{Z}(d|\mathcal{M}) = \int p(\theta|d, \mathcal{M})\pi(\theta|\mathcal{M})d\theta, \quad (2)$$

which describes the success at which the hypothesis represents the data. Parameter estimation pipelines such as *Bilby* assume GW noise data to be stationary and Gaussian [4]. The likelihood for transient behaviours present in GW strain data is thus expressed using the following Gaussian noise likelihood \mathcal{L} , with a data value d_k at a frequency bin index k [3, 7]:

$$\ln \mathcal{L}(d|\theta) = -\frac{1}{2} \sum_k \left\{ \frac{[d_k - \mu_k(\theta)]^2}{\sigma_k^2} + \ln(2\pi\sigma_k^2) \right\}, \quad (3)$$

where σ_k is the amplitude spectral density for the noise at a given frequency bin and $\mu_k(\theta)$ is the waveform in that frequency bin. The non-Gaussianity of transient glitches contradicts this assumption, further demonstrating the importance of producing a means to remove these triggers from GW data.

II. OBJECTIVE

We have constructed a parametric model which provides a baseline to identify fast scattering glitches from GW data and mitigate these instances for improved analysis of CBC events. Parameterised models provide a more reliable method of both the subtraction and marginalisation of long-duration glitches as opposed to *BayesWave* [9, 10], which is more reliable for its wavelet models and greater proficiency in subtracting short-duration glitches. Parametric model inference, performed primarily using *Bilby*, provides a more robust probe of glitch morphology by leveraging data from the entire segment duration to influence our inference of the glitch morphology at specific time/frequency zones overlapping a CBC [3]. The construction of this model enables us to simulate the characteristics and conditions of fast scattering glitches. By inferring their parameters, we can assess the likelihood that specific configurations may approximate a fast scattering glitch observed in GW strain data.

III. METHODS

A. Model Construction

In constructing a mathematical model for a generic fast scattering glitch waveform, we began by following the form of the undermentioned equation that describes the excess strain noise $h(t)$ related to the motion of the surface $x(t)$ produced by scattered light of wavelength λ over time t [2-4, 6]:

$$h(t) = \bar{A} \sin \left[\frac{4\pi}{\lambda} x(t) + \phi \right], \quad (4)$$

where \bar{A} is the amplitude of the noise produced by the glitch with a phase shift ϕ .

Assuming the movement of the relevant surface as a harmonic oscillator, its motion may be presented as such with the incorporation of the two driving frequencies f_1 and f_2 , the two of which interact with one another to produce fast scattering glitches which are then seen in GW strain data:

$$x(t) = A_1 \sin(2\pi f_1 t) + A_2 \sin(2\pi f_2 t), \quad (5)$$

where A_1 and A_2 are the respective amplitudes associated with the two driving frequencies.

Incorporating the above equation into Equation 4, we thus arrive at our model for a fast scattering glitch:

$$h(t) = \bar{A} \sin \left[A_1 \sin(2\pi f_1 t) + A_2 \sin(2\pi f_2 t) + \phi \right]. \quad (6)$$

B. Model Testing: Preliminaries

1. Spectrogram Predictions

Testing the validity of the model began by determining general predictions for each driving frequency and their associated amplitudes by incorporating the following relation between the wavelength of the scattered light beam λ , its frequency f , and velocity of propagation $v(t)$ [3, 6]:

$$f = \left| 2 \frac{v(t)}{\lambda} \right|. \quad (7)$$

We thus determine the velocity associated with the motion from Equation 5, and by simplifying for the case of the maximum frequency f_{max} where $\cos(2\pi ft) = 1$:

$$A_1 f_1 + A_2 f_2 = f_{max}. \quad (8)$$

Because fast scattering glitches are commonly characterised by a peak frequency of 50 Hz as seen in detector data, we assume that $f_{max} = 50$ when determining the possible values of each parameter associated

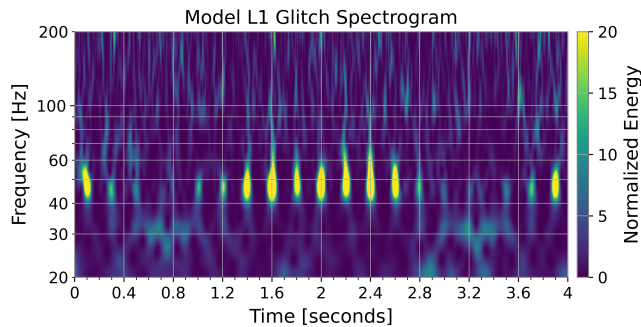


FIG. 2: A spectrogram produced by the mathematically derived fast scattering model with parameters $f_1 = 0.2$, $f_2 = 5$, $d_1 = 35$, $d_2 = 15$, and $\phi = 0$, where $d_1 = \frac{A_1}{f_1}$ and

$$d_2 = \frac{A_2}{f_2}.$$

with Equation 8. We defined a list of prior assumptions for each parameter, drawn from uniform distributions: the prior ranges for d_1 and d_2 were inferred based on $f_1 \sim \mathcal{U}(0.1, 0.3)$ and $f_2 \sim \mathcal{U}(1, 5)$ via Equation 8, with $\bar{A} \sim \mathcal{U}(10^{-23}, 10^{-18})$ based on the characteristic LIGO strain and $\phi \sim \mathcal{U}(0, 2\pi)$. Figure 2 provides the resulting spectrogram for one of the possible parameter combinations of our fast scattering model, reproducing the long-duration, short-burst arches characteristic of traditional fast scattering glitch cases.

2. Parameter Injections & Waveform Generator

The next step in testing the efficacy of our model in fast scattering glitch emulation and mitigation involved determining the accuracy at which various given parameter injections align with the posterior results of each case. This was performed by running an analysis of the model with our chosen injection values through Bilby given the logical priors for each parameter. The prior assumptions include those of the ranges of the two driving frequencies characteristic of fast scattering triggers as mentioned previously, the possible values of d_1 and d_2 given the driving frequencies and f_{max} , the possible range of values associated with ϕ , and the noise amplitude \bar{A} .

Figure 3 displays a corner plot associated with a given set of parameter injection values, shown as singular points or lines overlaying a column or row of images. Corner plots are interpreted as follows: the outermost diagonal showcases one-dimensional histograms detailing the distribution of each parameter's values, while the lower triangle displays two-dimensional joint histograms representing two-parameter slices within the parameter space. Each distribution represents the a posteriori probability of the parameter configuration given the data, including the distribution for each parameter which, for this example, manifested as Gaussian probability curves.

Further testing of the validity of the glitch model involved generating a spectrogram from the injected parameter values into the model produced by Bilby's wave-




FIG. 3: A corner plot produced from Bilby given a set of injection parameters specified from Figure 2. The parameters d_1 and d_2 describe the amplitudes associated with the driving frequencies f_1 and f_2 .

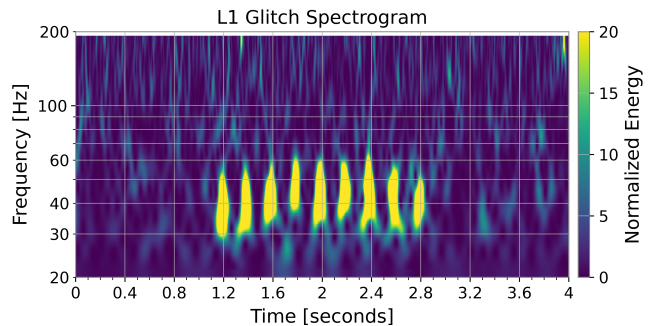


FIG. 4: The resulting spectrogram from the Bilby waveform generator given the previously specified set of injection parameters in Figures 2 and 3.

form generator. The shape and pattern of the resulting spectrogram fell in close agreement with the spectrogram produced directly from the mathematical model, as shown in Figure 4.

3. P-P Testing

The final validity test which we performed consisted of P-P testing, or parameter-parameter testing, to ensure unbiased posterior results given a set of samples taken from the prior. P-P plots compare the empirical cumulative distribution function (CDF) versus the theoretical CDF. In a P-P plot, an unbiased result in which the data's distribution matches the theoretical distribution




FIG. 5: The P-P plot produced for the model with a sample size of $N = 100$. The three shaded error regions represent credible bands; the darkest shade signifies the highest credibility, tapering off to lower levels of credibility in lighter-shaded areas.

will manifest as a diagonal plot where 10% of the samples have $CDF < 0.1$, 20% have $CDF < 0.2$, and so on. Bias, however, will manifest as a deviation from the diagonal. Figure 5 displays the associated P-P plot for a sample size of $N = 100$ simulated signals, where each parameter curve is shown to lie within the 3σ uncertainty. According to our P-P plot, the events within each credible interval (CI) delineate a clear linear progression and thereby suggest unbiased parameters.

C. Model Testing: Glitch Mitigation

The culmination of the findings detailed in Sections III B 1, III B 2, and III B 3 demonstrated our model’s potential in accurately depicting true fast scattering triggers. We thus moved on to our next phase of model testing, which was done on two different manifestations of fast scattering events: hardware-injected triggers and real triggers, the results of which are detailed in Sections III C 1 and III C 3, respectively. The former was to determine our model’s efficiency in glitch emulation and mitigation for a simpler fast scattering structure case, while we intended the latter to serve as a step up in difficulty for more complex fast scattering structures.

1. Hardware-Injected Trigger Events

Hardware injections often serve to simulate GW signals using repeated injections in order to calibrate the detectors and to provide a test of parameter estimation analyses in recovering signals from the data collected by the detectors. Additionally, hardware injections may serve to imitate non-Gaussian noise, including fast scattering

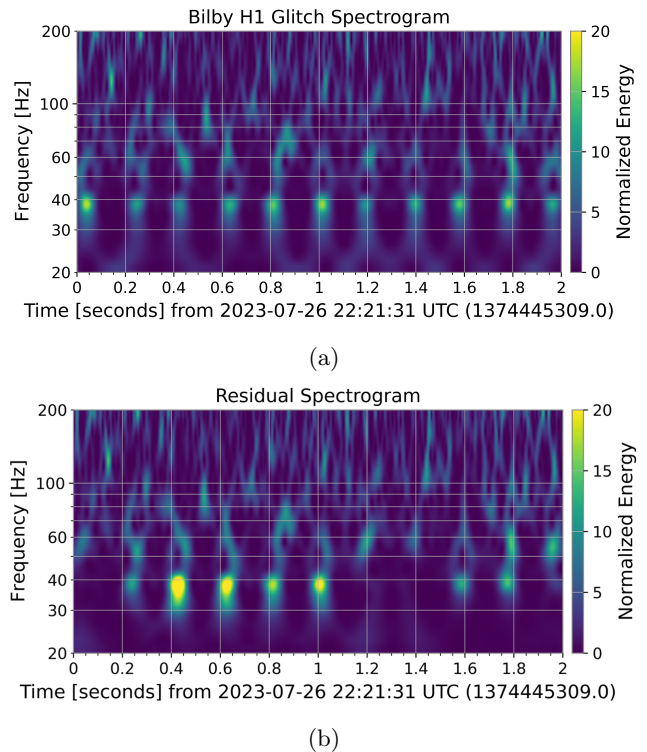


FIG. 6: The data spectrogram (6a) and its residual spectrogram (6b) of a hardware-injected fast scattering glitch administered at the Hanford site. In the residual, the model effectively mitigated some frequency peaks but amplified others.

triggers. Hardware injections are performed by physically displacing the test masses within the detector using an actuator in order to replicate the movements that occur as a result of GW signals or environmental oscillations [11].

The spectrogram for the case of interest is given in Figure 6a, in which a 2.6 Hz sine wave was injected in the detector. We performed parameter estimation on the hardware-injected event and recovered a posterior probability distribution, which then yielded the maximum likelihood values for each parameter. Presented in Figure 7 is the resulting corner plot. As expected, we recovered the known driving frequency $f_2 = 2.6$ Hz, which indicated that Bilby produced a favourable result in its parameter estimation. Figure 6b displays the residual spectrogram after our glitch mitigation efforts. It is evident that the model’s success was limited: triggers at $t = 0$ s, $t = 1.2$ s, $t = 1.4$ s, and $t = 2$ s were removed, while triggers from $t = 0.2 - 1$ s and $t = 1.6 - 1.8$ s persisted; in some cases, they were amplified, which we attribute to possible injection amplitude changes.

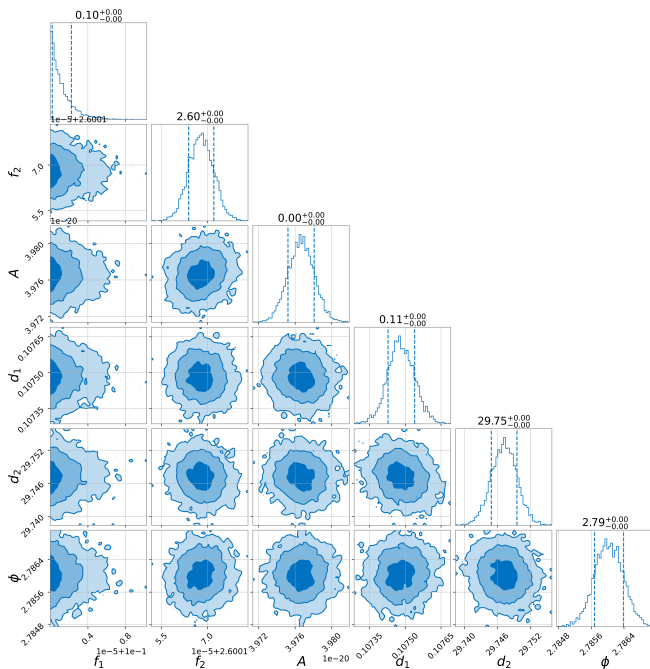


FIG. 7: The corner plot providing the posterior parameter distributions of the hardware-injected event. Our posterior analysis accurately reflected a value of $f_2 = 2.6$ Hz characteristic of the shaker injection applied at the LIGO Hanford detector.

2. Real Trigger Events: O3 Example

We adopted the fast scattering glitches seen in Figure 1 as provided by [2] from LIGO’s third observing run to perform parameter estimation analyses on a fast scattering instance of a more complex structure.

In Figure 8, we provide a spectrogram which displays a fast scattering glitch arch and its respective residual. The results implied that the model was less effective in mitigating the fast scattering triggers analysed: the residual, as opposed to any mitigation, instead indicates an addition to data.

3. Real Trigger Events: GW190701

We tested our model’s mitigation capabilities on an instance of fast scattering which persisted in the presence of a CBC merger, an event known as GW190701. We performed the first investigation in our model’s abilities to perform fast scattering mitigation in the presence of GW190701 for a fast scattering arch that manifested prior to the CBC merger. Figure 9 displays the data spectrogram of the with its associated residual. This arch is characterised by irregular peaks which persist at varying time intervals. The asymmetry of this pattern hindered our model’s effectiveness, a shortcoming evident from the minimal mitigation portrayed in the residual.

Upon performing parameter estimation on the fast

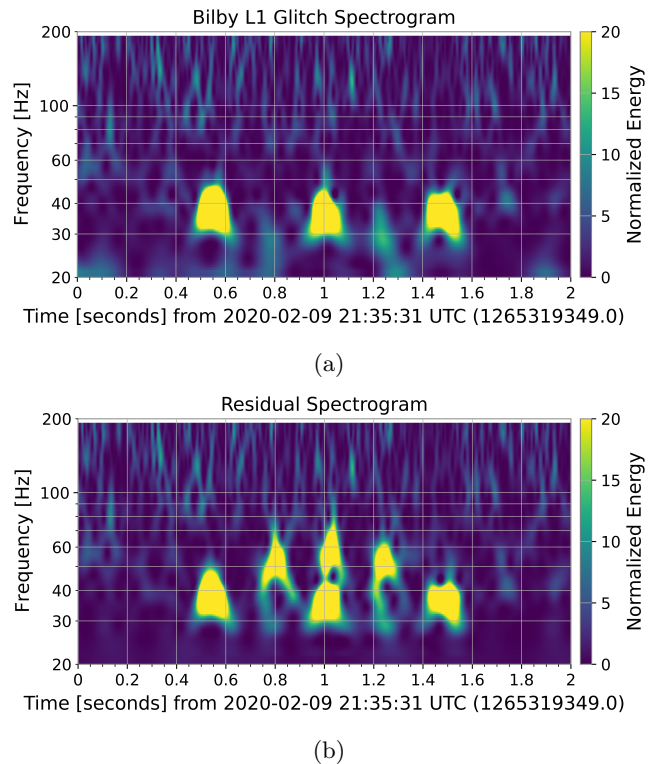


FIG. 8: The spectrogram (8a) and its residual (8b) of a fast scattering arch as seen from [2]. Contrary to glitch mitigation efforts, the residual reveals newly introduced fast scattering peaks.

scattering structure superimposed on the merger event, our model demonstrated remarkable proficiency. It effectively replicated this structure, eliminating the offending glitch as seen in Figure 10. This outcome not only indicates our model’s potential in performing robust glitch mitigations but also showcases an enhancement over mitigation efforts by Abbott et. al. (2021) [12] with *BayesWave*, as illustrated in Figure 11.

IV. CONCLUSIONS

We constructed a parametric model for fast scattering glitches under the assumption that the motion of the reflecting instrumentation inside the LIGO detectors adheres to the behaviour of a simple harmonic oscillator. After developing a reliable model, we performed a variety of validity tests, including generating spectrogram predictions to describe the general waveform of a fast scattering glitch, performing parameter estimations with injections, and P-P testing. After our model passed these preliminary assessments, we graduated to parameter estimations and mitigations of actual glitch occurrences seen in GW strain data. This included a hardware-injected trigger event and a naturally occurring fast scattering event. We then executed the same analyses on a fast scattering arch preceding GW190701 and another over-

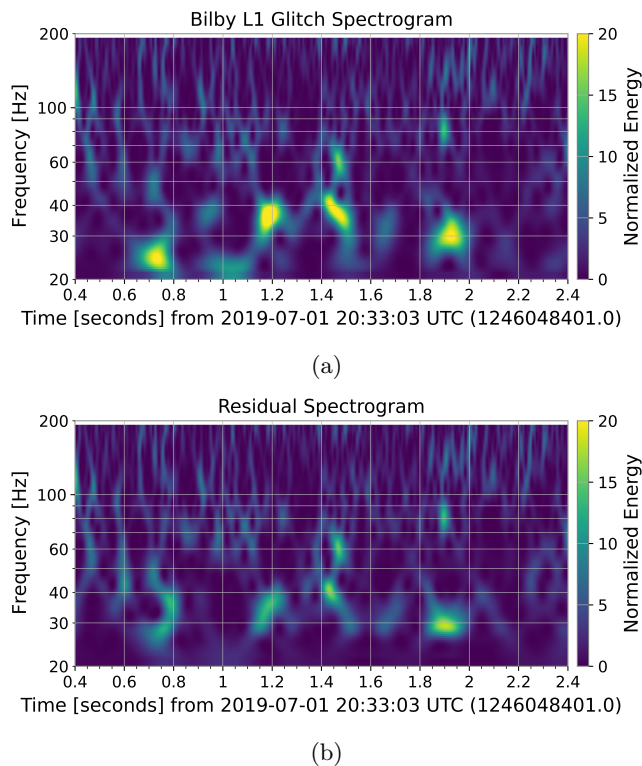


FIG. 9: The data spectrogram (9a) and its residual (9b) of the fast scattering triggers existing before GW190701. The residual spectrogram shows minimal mitigation of individual triggers, marking the difficulty our model encountered when attempting to replicate intricate glitch structures.

lapping the merger.

Though our parametric model faced challenges in accurately imitating the fast scattering cases in Sections III C 1 and III C 2, we achieved notable progress with enhanced glitch mitigation for GW190701, detailed in Section III C 3, surpassing previous efforts as performed in Abbott. et. al. (2021) [12]. This outcome accentuates the formidable potential of our parameterised model in adeptly mitigating fast scattering glitches. To further hone its performance, we are committed to refining the model in future work. Our intended approach will involve testing it against a broader array of fast scattering instances in LIGO detector data, continuing to fine-tune its parameters accordingly, and conducting parameter estimations on joint “CBC+glitch” examples. These objectives pave the way towards our primary aim: enhancing the analysis of GW events and bolstering signal detection in forthcoming observing runs.

V. ACKNOWLEDGEMENTS

I am immensely thankful to my mentors Rhiannon Udall and Derek Davis for their guidance, support, understanding, and constructive feedback. I also thank the

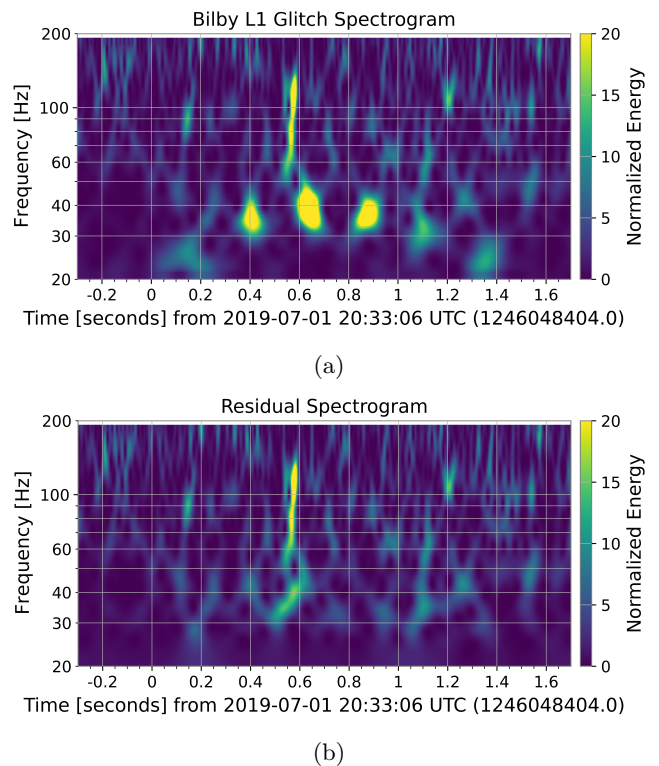


FIG. 10: The data spectrogram (10a) and its residual (10b) of the fast scattering segment overlapping GW190701. The persisting fast scattering frequency peaks are notably completely absent post-mitigation.

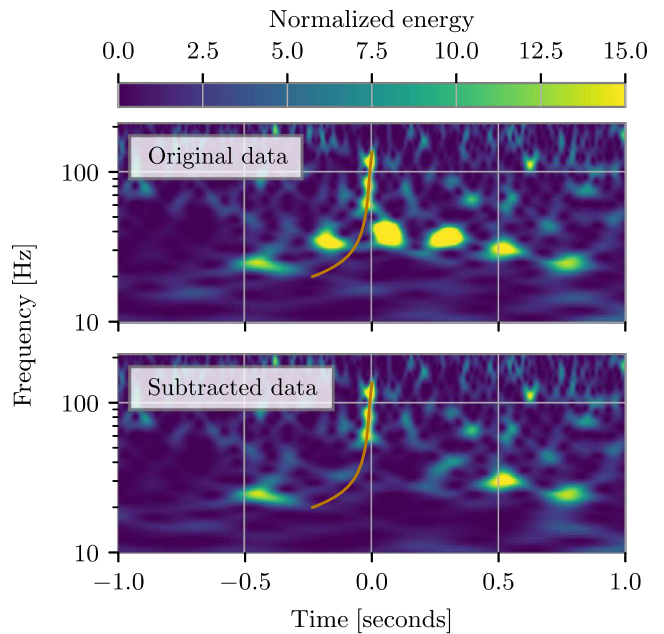


FIG. 11: The glitch mitigation performed using BayesWave in [12] on top of GW190701, represented as a red line in each spectrogram. The top spectrogram represents the data pre-mitigation, while presented below is the residual post-mitigation efforts.

LIGO Laboratory group at the California Institute of Technology for their advice and for fostering a welcoming and positive work environment. This work was graciously made possible through funding provided by the National Science Foundation. I gratefully acknowledge

the Student-Faculty Program and the LIGO Summer Undergraduate Research Fellowship Program at the California Institute of Technology for the rewarding summer experience. Finally, I thank the Cal-Bridge Program for their encouragement and support.

-
- [1] Benjamin P Abbott, R Abbott, TD Abbott, MR Abernathy, F Acernese, K Ackley, C Adams, T Adams, P Addesso, RX Adhikari, et al. Gw150914: The advanced ligo detectors in the era of first discoveries. *Physical review letters*, 116(13):131103, 2016.
- [2] S Soni, C Austin, A Effler, RMS Schofield, G González, VV Frolov, JC Driggers, A Pele, AL Urban, G Valdes, et al. Reducing scattered light in ligo’s third observing run. *Classical and Quantum Gravity*, 38(2):025016, 2020.
- [3] RP Udall and Derek Davis. Bayesian modeling of scattered light in the ligo interferometers. *Applied Physics Letters*, 122(9), 2023.
- [4] Arthur E Tolley, Gareth S Cabourn Davies, Ian W Harry, and Andrew P Lundgren. Archenemy: removing scattered-light glitches from gravitational wave data. *Classical and Quantum Gravity*, 40(16):165005, jul 2023.
- [5] Siddharth Soni, Christopher Philip Luke Berry, Scott B Coughlin, Mahboobeh Harandi, Corey B Jackson, Kevin Crowston, Carsten Østerlund, Oli Patane, Aggelos K Katsaggelos, Laura Trouille, et al. Discovering features in gravitational-wave data through detector characterization, citizen science and machine learning. *Classical and Quantum Gravity*, 38(19):195016, 2021.
- [6] T Accadia, F Acernese, F Antonucci, P Astone, G Ballardini, F Barone, M Barsuglia, Th S Bauer, MG Beker, A Belletoile, et al. Noise from scattered light in virgo’s second science run data. *Classical and Quantum Gravity*, 27(19):194011, 2010.
- [7] Gregory Ashton, Moritz Hübner, Paul D Lasky, Colm Talbot, Kendall Ackley, Sylvia Biscoveanu, Qi Chu, Atul Divakarla, Paul J Easter, Boris Goncharov, et al. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(2):27, 2019.
- [8] Isobel M Romero-Shaw, C Talbot, S Biscoveanu, V D’emilio, G Ashton, CPL Berry, S Coughlin, S Galadage, C Hoy, M Hübner, et al. Bayesian inference for compact binary coalescences with bilby: validation and application to the first ligo–virgo gravitational-wave transient catalogue. *Monthly Notices of the Royal Astronomical Society*, 499(3):3295–3319, 2020.
- [9] Neil J Cornish and Tyson B Littenberg. Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches. *Classical and Quantum Gravity*, 32(13):135012, jun 2015.
- [10] Sophie Hourihane, Katerina Chatzioannou, Marcella Wijngaarden, Derek Davis, Tyson Littenberg, and Neil Cornish. Accurate modeling and mitigation of overlapping signals and glitches in gravitational-wave data. *Phys. Rev. D*, 106:042006, Aug 2022.
- [11] C Biwer, D Barker, JC Batch, J Betzwieser, RP Fisher, Evan Goetz, Shivaraj Kandhasamy, S Karki, JS Kissel, Andrew P Lundgren, et al. Validating gravitational-wave detections: The advanced ligo hardware injection system. *Physical Review D*, 95(6):062002, 2017.
- [12] R Abbott, TD Abbott, S Abraham, F Acernese, K Ackley, A Adams, C Adams, RX Adhikari, VB Adya, Christoph Affeldt, et al. Gwtc-2: compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *Physical Review X*, 11(2):021053, 2021.