



GRAVITATIONAL WAVEFORM MODELLING FOR INFERENCE OF SPIN PRECESSION IN BINARY BLACK HOLE MERGERS

By

LUCY MAJA THOMAS

A thesis submitted to
the University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Astrophysics and Space Research Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
June 2023

RESEARCH THESIS *Author's Declaration*

Full name (block capitals, surname first): LUCY MAJA THOMAS


Full title of thesis/dissertation (block capitals): GRAVITATIONAL WAVEFORM MODELLING

..... FOR INFERENCE OF SPIN PRECESSION IN BINARY BLACK HOLE MERGERS

College/School/Department (block capitals): SCHOOL OF PHYSICS AND ASTRONOMY

Date of award of degree (leave blank):

1. I understand that one printed and one electronic copy of my thesis/dissertation (the Work) will be deposited in the University Library (the Library) and in a suitable electronic repository, for permanent retention.
2. Without changing the content, the Library or any third party with whom it has an agreement to do so, may convert either copy into a format or medium for the purpose of long-term retention, accessibility and preservation.
3. The Library will publish, and/or arrange with appropriate third parties for the non-exclusive publication of, a bibliographic description of the thesis/dissertation, and the author's abstract.
4. Unless arrangements are made to the contrary, (see paragraph 6. below), the Library is authorised to make the Work available for consultation in the Library, and via a recognised inter library loans system. The Library is authorised to make the electronic copy of the Work freely accessible to individuals and institutions - including automated agents - via the Internet.
5. Rights granted to the University of Birmingham through this agreement are entirely non-exclusive. I retain all my rights in the Work in its present version or future derivative works.
6. I understand that I may apply to the University to retain the right to withhold access to the content of my thesis/dissertation. Access to the paper version may be withheld for a period which shall not normally exceed four calendar years from the congregation at which the degree is conferred. The actual length of the period will be specified in the application, together with the precise reasons for making that application. The electronic copy may be withheld from dissemination via the web or other networks for any period.
7. I have obtained permission for any use made of substantial amounts of published or unpublished copyright material (text, illustrations, etc) where the rights are owned by a third party, other than as permitted under either The Copyright Designs and Patents Act 1988 (as modified by any related or successor legislation) or the Terms and Conditions of any Licence governing its use.
8. The content of the copies I shall deposit with the Library will be the final version of my thesis, as approved by the Examiners.
9. I understand that the Library and administrators of any electronic theses repository do not hold any obligation to take legal action on behalf of myself, or other rights holders, in the event of a breach of intellectual property rights, or any other right, in the material deposited.
10. I understand that, in the event of my thesis/dissertation being not approved by the Examiners, this declaration will become null and void.

Signature: 

Date: 3rd October 2023

For Library use (please leave blank):

Classmark:

Accession number:

Control number:

eTheses Repository url:




Figure 1: A quasicircular inspiralling binary black hole, drawn by L M Thomas.

ABSTRACT

Detections of gravitational waves from the mergers of compact binaries have revolutionised our understanding of the universe. As well as confirming a key prediction of Einstein's Theory of General Relativity, they have allowed unique opportunities to probe strong gravity, observe previously dark populations of binary black holes, and constrain binary formation channels.

Spin precession occurs due to the coupling of the spins of individual black holes with their orbital motion, resulting in complex and rich dynamics that encode valuable astrophysical information. To date no confident detection of a precessing binary merger has been made, in part due to systematics in the waveform models used to infer the existence of precession. However precession has been shown to exist at the level of the population, making individual detections a tantalising possibility.

This thesis addresses two key questions in waveform modelling of precession. Firstly, we develop a strategy to reduce the high dimensionality of precessing binaries, which may be crucial for ensuring accuracy of precessing waveform models through calibration to numerical relativity. Secondly, we utilise the power of artificial neural networks to build a waveform model that faithfully mimics a highly accurate precessing multipolar waveform, with a fraction of the evaluation cost. Since the detection of precession is predicated on waveform models which accurately contain the relevant physics and can be practically used for inference, striking this delicate balance between efficiency and accuracy is vital.

Finally, we consider the impact of low frequency sensitivity upon detections of precession in a black hole binary population. Next-generation ground-based detectors are expected to improve upon the low frequency sensitivity of current instruments, providing access to more inspiral content in binary mergers. We show that this will allow us to converge upon the true underlying population spin distribution faster, which could lead to more accurate constraints upon the formation pathways of binary black holes.

To my Mum, without your immeasurable love and support this would not have been possible.

To my Dad, your immense pride in me helped me believe that it was possible.

ACKNOWLEDGMENTS

My heartfelt thanks go to my fantastic supervisor, Patricia. Your support and mentorship have been invaluable to me these four years, and have shaped me into the scientist I am today. I also thank Geraint, whose patience, insight and guidance have always kept me on track. Without the two of you, there would be no thesis and I am immensely grateful to you both.

My thanks also go to my other colleagues within the ASR group at Birmingham, and to the people I have been lucky enough to work with in the LVK collaboration. It's been a thrill to see such cutting-edge science happening all around me and be a small part of it. Thank you to my thesis examiners, Mark and Guy, for fruitful discussion and a surprisingly enjoyable viva examination. Also to Annelies, for your support and for making labs fun.

My thanks go to my fellow PhD students and friends at Birmingham, for the Monday pub quizzes, Thursday pub clubs, and countless coffee breaks which sometimes lasted hours: Alex, Daria, Diganta, Eliot, Emily, Evan, Hannah, Lucy, Martin, Matt, Owen, Piero, Ruairi, Shola, Xinyue, and Yasmin. Thank you especially to Natalie and Alice, for always being on hand to ask and answer stupid questions and listen to me complain when my code was broken; maybe this thesis can serve as an example of what not to do. To Alex 'Jill' and to Aysha, for making me laugh even on the most stressful days and helping me to procrastinate. To Cress and Dan, your friendship has been a constant through the best and worst of this PhD. Thank you to Cress for always being there with wine and a chat when I needed, and to Dan, whose terrible jokes made me realise that it could always be worse. Thank you to Thomas, for your unfailing love and support, and meticulous proofreading.

Finally, an enormous thank you to my family. To Ben, Charlotte, and little Sophie, to Omi and Chris, and especially to my parents. You were only ever a phone call away

which means the world to me. To Elsie and Dottie, for the cuddles whenever I needed them.

This thesis and the work I undertook as part of my PhD is the result of an estimated 270,000 CPU hours of computing time. This equates to roughly 13.2 tonnes of CO₂ [1], or the equivalent of 8 flights between London and Los Angeles.

Contents

	Page
1 Introduction	1
2 Gravitational Waves: Generation and Detection	5
2.1 Introduction to Gravitational Waves in General Relativity	5
2.1.1 Gravitational Waves in Linearised Gravity	6
2.1.2 Sourcing the Gravitational Waves	10
2.2 Astrophysical Sources of Gravitational Waves	13
2.2.1 Compact Binary Coalescences	13
2.2.2 Detecting Gravitational Waves from Black Hole Mergers	19
2.2.3 Other Astrophysical Sources of Gravitational Waves	25
2.3 Modelling Gravitational Waveforms	26
2.3.1 Weak Field Solutions and the Post-Newtonian Approximation	26
2.3.2 Numerical Relativity and the Strong Field Regime	30
2.4 Effects of Black Hole Spin and Precession	33
2.4.1 Aligned Spin Binaries	33
2.4.2 Spin Precession	35
2.5 Astrophysical Inference from Gravitational Wave Detections	37
2.6 Summary	41
3 Waveform Modelling for Scientific Interpretation of Gravitational Waves	43
3.1 Multipole Waveform Decomposition	44
3.2 Modelling Precession	49

3.3	Effective-One-Body Models	54
3.3.1	Effective-One-Body Hamiltonian	55
3.3.2	Gravitational Waveform Modes	56
3.3.3	Radiation-Reaction Force	56
3.3.4	Full Formalism	57
3.4	Phenomenological Models	58
3.4.1	Closed-Form Expressions and Phenomenological Fits	59
3.4.2	Modelling Across Parameter Space	61
3.4.3	Full Inspiral-Merger-Ringdown Waveforms	62
3.4.4	Including Precession in Phenom Waveforms	62
3.5	Surrogate Models	64
3.6	Summary	67
4	A New Effective Precession Spin for Modelling Multimodal Gravitational Waveforms in the Strong-Field Regime	69
4.1	Calibrating Waveform Models to Precessing Numerical Relativity	70
4.1.1	A New Effective Precession Spin	75
4.2	Methodology	81
4.2.1	Waveforms	81
4.2.2	Faithfulness for Precessing Waveforms	82
4.2.3	Binary Configurations	85
4.3	Results	88
4.3.1	Mode Analysis	88
4.3.2	Strain Analysis	92
4.3.3	Accuracy of the Final Spin and Recoil	99
4.4	Discussion	102
5	Accelerating Multimodal Gravitational Waveform Models from Precessing Compact Binaries with Artificial Neural Networks	109

5.1	Methodology	111
5.1.1	Artificial Neural Networks	111
5.1.2	Waveform Decomposition	113
5.2	Model	115
5.2.1	Training, Validation and Testing Data	115
5.2.2	Reduced Basis and Empirical Interpolant	119
5.2.3	Parameter Space Fits with Artificial Neural Networks	122
5.2.4	Complete Surrogate Model	127
5.3	Model Evaluation	133
5.3.1	Waveform Accuracy	133
5.3.2	Timing	136
5.4	Discussion	139
6	Constraining Precession in the Black Hole Binary Population with Next-Generation Ground-Based Gravitational Wave Detectors	143
6.1	Hierarchical Bayesian Analysis	147
6.1.1	Hypermodel of Spin Magnitude Distribution	150
6.1.2	Hypermodel of Effective Spin Parameter Distribution	151
6.2	Systematic Injection Study	152
6.2.1	Methodology	152
6.2.2	Results	158
6.3	Population Study	161
6.3.1	Population Binary Parameters	162
6.3.2	Selected Individual Binary Parameter Estimation Results	170
6.3.3	Population Inference Results	179
6.4	Discussion	189
7	Conclusion	193

List of Figures

1	A quasicircular inspiralling binary black hole, drawn by L M Thomas.	2
2.1	A binary black hole moving on a circular orbit.	14
2.2	Polarisations h_+ , h_\times of a binary black hole coalescence showing the inspiral, merger and ringdown.	18
2.3	Black hole binary with orientation relative to an observer.	20
2.4	Polarisations h_+ , h_\times of a precessing black hole binary coalescence.	37
3.1	Amplitudes in the frequency domain for the spherical harmonic modes $\tilde{h}_{\ell,m}(f)$ for a non-spinning black hole binary.	47
3.2	The posterior distribution for the luminosity distance d_L and inclination θ_{JN} for the gravitational wave event GW190412.	49
3.3	Definition of the inertial J -frame and the Euler angles.	53
4.1	Distribution of the 652 NR simulations used in calibration of PhenomXHM	72
4.2	Amplitude of the $(2, 1)$ -mode (left) and the $(2, -1)$ -mode (right) for a fiducial precessing binary, showing the fully precessing waveform and waveforms parameterised by χ_p and χ_\perp	77
4.3	Amplitude of the $(2, 1)$ -mode, and time evolution of two quaternion components, for a fiducial binary. We show the fully precessing components, and those parameterised by χ_p and χ_\perp	78
4.4	Amplitude of the $(2, 1)$ -mode, and time evolution of two quaternion components, for an equal mass fiducial binary. We show the fully precessing components, and those parameterised by χ_p and χ_\perp	79

4.5	Amplitude of the (2, 1)-mode, and time evolution of two quaternion components, for a fiducial binary with mass ratio 3. We show the fully precessing components, and those parameterised by χ_p and χ_\perp	80
4.6	Cumulative histograms of white noise matches for the (2, 2)-mode and the (2, 1)-mode.	89
4.7	Complete results for white noise matches between the fully spinning waveform and each of the effective spin parameterisations, $\vec{\chi}_\perp$ and χ_p	90
4.8	Full results for the O4 PSD-weighted matches between the fully spinning waveform and each of the effective spin parameterisations, $\vec{\chi}_\perp$ and χ_p	93
4.9	Amplitude of the waveform strain $h(t)$ for the same fiducial binary as in Figure 4.2 at an inclination of $\iota = \pi/3$, showing the fully precessing waveform and the two effective spin parameterisations, χ_p and $\vec{\chi}_\perp$	94
4.10	Histograms of the sky-and-polarization-averaged strain mismatches $\overline{\mathcal{M}\mathcal{M}}_{\text{strain}}$ between the fully precessing waveform and each of the two-spin mappings using the O4 PSD.	95
4.11	SNR-weighted strain mismatches $\overline{\mathcal{M}\mathcal{M}}_{\text{SNR}}$ as a function of binary total mass M for 100 binaries for the $\vec{\chi}_\perp$ and χ_p mappings.	96
4.12	Sky-and-polarization-averaged strain mismatch $\overline{\mathcal{M}\mathcal{M}}_{\text{strain}}$ versus $\Delta\lambda_L(t_0)$ for the 20,833 binaries with both the $\vec{\chi}_\perp$ and χ_p effective spin parameterisations.	97
4.13	Time evolution of total in-plane spin magnitude S_\perp in the coprecessing frame for the same fiducial binary as in Figure 4.2, with the fully precessing configuration shown alongside the two effective spin mappings.	97
4.14	Time evolution of the quaternion components \hat{q}_0 and \hat{q}_3 , and \hat{q}_1 and \hat{q}_2 for the fiducial precessing binary, for the fully precessing configuration as well as the two effective spin mappings.	98

LIST OF FIGURES

4.15 Cumulative distribution of matches for two of the four quaternion elements \hat{q}_1 and \hat{q}_2 , between the fully precessing dynamics and each of the $\vec{\chi}_\perp$ and χ_p -mapped systems. 99

4.16 Error in the azimuthal angle of the final spin state, $\Delta\phi_f$, in radians, between the final spin state produced by the fully precessing waveform, χ_f , and the resulting final spin state of the waveform produced by the $\vec{\chi}_\perp$ and χ_p mappings. 101

4.17 Absolute value of the error in the recoil velocity magnitude Δv_f , and recoil velocity tilt angle $\Delta\theta_{v_f}$, between the fully spinning waveform and each of the $\vec{\chi}_\perp$ - (teal) χ_p - mappings. 103

4.18 Cumulative histograms of white noise mode-by-mode matches for the (2,2)-mode and the (2,1)-mode for mass ratios $q = 1$ and $q = 3$ for the same binaries as in Figure 4.6 with the $\vec{\chi}_\perp$ -parameterisation and the χ_p -parameterisation. 107

5.1 Visualisation of the spin parameters of the entire training dataset coloured by number density. 116

5.2 Graphical representation of the ANN architecture for the coprecessing (2, 2)-mode phase $\phi_{22}(t; \vec{\lambda})$, as an example. 124

5.3 Training and validation losses for the (2, 2)- and (2, 1)-mode, for both amplitude and phase. 124

5.4 Training and validation losses for the (3, 3)- and (4, 4)-modes, both amplitude and phase. 126

5.5 Comparison of the coprecessing mode amplitudes and phases, Euler angles, and time-domain strain, between the neural network-predicted waveform components and the true SEOBNRv4PHM data, for a fiducial precessing binary. 129

5.6 White noise mismatches between the SEOBNRv4PHM -generated coprecessing frame mode data and and the neural network-predicted coprecessing mode, for each of the four modes across the 10, 000 binary test set. 130

5.7	White noise mismatches between the SEOBNRv4PHM-generated coprecessing frame mode data and the neural network-predicted coprecessing mode, for the (2, 2) and the (2, 1)-mode, as computed in Fig. 5.6 but for a range of different total masses.	131
5.8	Worst 5% of test dataset mismatches for the coprecessing (2, 2)- (2, 1)-modes, shown in parameter space of mass ratio q against in-plane spin magnitude $ \chi_{1\perp} $	132
5.9	Worst 5% of test dataset mismatches $\bar{\mathcal{M}}_f$ for the (3, 3)- and (4, 4)-modes, shown in parameter space of mass ratio and in-plane spin magnitude, and coloured by mismatch.	133
5.10	Time-domain mismatches for the surrogate model for the Euler angles against the training data.	134
5.11	Orientation averaged mismatches for SEOBNRv4PHM against SEOBNN_v4PHM_4dq2 for all ≤ 4 modes in the J -frame.	136
5.12	Computational cost for each step in the waveform construction.	138
5.13	Computational cost per binary for evaluating the 22-mode surrogate model over batches of N_λ binaries.	139
6.1	PSDs for proposed next generation detectors Cosmic Explorer, the Einstein Telescope, as well as the proposed LIGO A+ upgrade, and the Virgo design sensitivity curve.	146
6.2	Posterior distributions for the spin parameters of the GW190521-like injected binary whose parameters are shown in Tab. 6.1, for the fixed d_L systematic series.	157
6.3	Posterior distributions for spin parameters of the GW190521-like injected binary whose parameters are shown in Tab. 6.1, for the series with fixed network ρ_{network}	159

6.4	Median PPDs, 90% credible intervals for the primary and secondary spin magnitudes distribution χ_1 , χ_2 , drawn from the inferred population from GWTC-3. We also plot the best-fitting Beta distributions to the median PPDs, and the Beta distribution which best fits our sample population of 20 binaries.	164
6.5	Median PPDs, 90% credible intervals for the effective spin parameter distributions χ_{eff} , χ_p , drawn from the inferred population from GWTC-3. We also plot the best-fitting 2D Gaussian to the underlying population, and the 2D Gaussian which best fits our sample population of 20 binaries. . . .	165
6.6	Distribution of the detector frame total mass M in solar masses M_\odot and χ_p for the 20 binaries selected from the population, coloured by the inclination angle θ_{JN}	168
6.7	Number of precession cycles $\alpha(f_{\text{low}}, f_{\text{ISCO}})/2\pi$ for the 20 selected binaries. .	169
6.8	90% 2D credible intervals and 1D posterior distributions for the detector-frame mass parameters of representative binary 6 for each of the 3 f_{low} values.	173
6.9	Parameter estimation results for spin parameters for representative binary 6 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$, 10Hz , and 5Hz , and plot the true injected values as shown in Tab. 6.2.	174
6.10	90% 2D credible intervals and 1D posterior distributions for the detector-frame mass parameters of binary 5 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$, 10Hz , and 5Hz , and plot the true injected values as shown in Tab. 6.2.	175
6.11	Parameter estimation results for spin parameters for binary 5 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$, 10Hz , and 5Hz , and plot the true injected values as shown in Tab. 6.2.	176
6.12	$\Delta\chi_{\text{eff}}$ and $\Delta\chi_p$ for the 20 selected binaries, defined as the deviation of the posterior samples from the true injected χ_{eff} and χ_p values for each binary. .	177

6.13	$\Delta\chi_1$ and $\Delta\chi_2$ for the 20 selected binaries, defined as the deviation of the posterior samples from the true injected χ_1 and χ_2 values for each binary. .	178
6.14	Results of population inference with the Gaussian Spin hypermodel. We show the posteriors on the hyperparameters $\{\mu_{\chi_{\text{eff}}}, \sigma_{\chi_{\text{eff}}}, \mu_{\chi_p}, \sigma_{\chi_p}, \rho\}$ for each of the three f_{low} values, 5Hz, 10Hz, and 20Hz.	180
6.15	Results of population inference with the Beta Magnitudes hypermodel. We show the posteriors on the hyperparameters $\{\alpha_{\chi_1}, \alpha_{\chi_2}, \beta_{\chi_1}, \beta_{\chi_2}\}$ for each of the three f_{low} values, 5Hz, 10Hz, and 20Hz.	184
6.16	Posteriors for spin magnitudes χ_1 and χ_2 . We show result for each of the three f_{low} values, showing individual event posteriors of spin magnitudes for the 20 binaries. We also show the Beta distributions of the median recovered posterior samples of $\{\alpha_{\chi_1}, \beta_{\chi_1}, \alpha_{\chi_2}, \beta_{\chi_2}\}$ for each value of f_{low} . Finally we show Beta distributions for the primary and spin magnitude with the underlying population hyperparameter values (column 3 of Tab. 6.4), and distributions with the injected hyperparameter values for the population of 20 binaries (column 4 of Tab. 6.4).	187

List of Tables

4.1	Mode energy thresholds for the odd m -modes.	84
4.2	Binary configurations used in the different match calculations.	87
5.1	Parameters of the 200,000 binaries which span our training dataset.	115
5.2	Greedy tolerances and reduced basis sizes for the amplitude and phase of each mode, as well as the Euler angles.	120
5.3	Maximum and median training and validation dataset mismatches for the amplitude and phase of each mode, as well as the Euler angles.	120
5.4	Details of the final neural network architecture for each component.	123
6.1	Injected and recovered parameters for the systematic series for a GW190521-like binary, including source frame masses, spin, and extrinsic parameters.	153
6.2	Injected intrinsic binary parameters, including detector frame masses, and spins, for the 20 binaries in our population.	166
6.3	Injected extrinsic binary parameters, including network SNR ρ_{network} for the 20 binaries in our population.	167
6.4	Injected and recovered values of the population hyperparameters for each of the Beta Magnitudes and Gaussian spin models, for each of the f_{low} cutoffs of 5Hz, 10Hz and 20Hz.	171

1. Introduction

Gravitational waves (GWs) are oscillating propagations of a gravitational field, just as electromagnetic waves such as gamma and radio waves are oscillating propagations of the electromagnetic field. Photons are emitted by accelerating charged particles, and analogously, accelerating masses produce GWs. GWs are a key prediction of Einstein's Theory of General Relativity (GR) [2] first published in 1916, and were first detected almost a century later in 2015 from the merger of a binary black hole (BBH), as a signal named GW150914 [254]. This landmark detection was a milestone in gravitational science, and there have now been around 90 detections of GW signals from mergers of binaries made up of black holes (BHs) and/or neutron stars (NSs) [3–5, 7, 45, 134, 218, 255–257] by the GW observatories Advanced LIGO [258, 259] and Virgo [243, 260]. These discoveries include the first multimessenger observation of a binary neutron star inspiral, GW170817 [261, 262], the first intermediate mass black hole GW190521 [37, 210], the first unequal-mass BBH GW190412 [109] and the first confident neutron star–black hole detections GW200105 and GW200115 [6]. Additional GW candidates have been reported from analyses of the publicly available data in Refs. [4, 7]. These compact binary merger observations have allowed us to probe gravity and test GR at a unique energy scale [8–12]. They have also had a transformative impact on our understanding of the properties of black holes and neutron stars, enabling us to constrain their mass and spin distributions [13] and to put them into their astrophysical context.

Astrophysical inference from gravitational wave data mainly relies on waveform models to extract source properties of the binaries, such as the masses and spins for BBHs. In order to fully exploit the scientific potential of these observations, waveform models need to be accurate, fast, and contain as much relevant physics as possible to accurately and efficiently extract the binary parameters. In the most generic BBHs, the spin vectors of the individual BHs may point in any direction. If the spins are aligned or anti-aligned with the orbital angular momentum of the binary, then the binary motion be restricted

to a fixed 2-dimensional plane known as the orbital plane. However, if one or both of the BH spins are misaligned with the orbital angular momentum, this will induce spin precessional motion of the orbital plane and BH spins [14, 15], inducing modulations to the GW amplitude and phase. This increased complexity in the signal makes precessing GWs more complicated to model, but the rich structure of precessing waveforms can allow us to break degeneracies in parameter space [16–19], and distinguish binary formation channels [20–36], so we can perform more reliable and complete inference. At the time of writing, there are a small number of detected binaries which may have been precessing [37, 38], and evidence of precession on the level of the BH population [34, 35], but both improved precessing waveform models and increased detector sensitivity may be required for more decisive measurements of precession.

The focus of this thesis is on modelling and inference of precession in BBHs with ground-based GW detectors. In Chapter 2 we introduce GWs as a natural consequence of GR, describing compact binary mergers as key sources of GWs for ground-based detectors and the key methods used to model them in the strong-field and weak-field gravity regimes. We then outline the phenomenology of these mergers, including the effects of spin precession, and give a brief overview of Bayesian inference for determining the binary source properties from GW detector data. In Chapter 3 we then introduce gravitational waveform modelling in more detail, introducing the three main approaches to waveform modelling with an emphasis on the inclusion of spin precession effects. In Chapter 4 we then move on to address a key problem in waveform modelling with precession: how to effectively incorporate precessing numerical information into semianalytic waveform models to increase their accuracy. This chapter is based on work published in Ref. [39], where we derive a new dimensional reduction strategy to accurately replicate precessing waveforms which include higher modes. In Chapter 5 we then focus on the complementary problem of ensuring that waveform models which include precession and higher modes are computationally efficient enough for practical use. This chapter is based upon Ref. [40], in which we build a surrogate waveform model using neural networks to speed up an underlying

model which includes precession and higher modes. In Chapter 6 we then address the question of how well we can expect to constrain precession in a population of BBHs with future ground-based GW detectors, whose low frequency sensitivity is as yet uncertain. Finally, in Chapter 7 we summarise and conclude this thesis.

2. Gravitational Waves: Generation and Detection

In Chapter 1, we introduced GWs as a propagation of a gravitational field and a prediction of GR. We explained that GW signals have been detected from mergers of compact binaries, and introduced the phenomenon of spin precession. In this Chapter, we motivate the study of GWs as a prediction of GR in more detail in Sec. 2.1, and then derive GWs in linearised gravity in Subsecs. 2.1.1 and 2.1.2. We then consider astrophysical sources of GWs in Sec. 2.2, including compact binary coalescences in Subsecs. 2.2.1 and 2.2.2, which form the main focus of this thesis, and other sources in Subsec. 2.2.3. We then discuss approximations to modelling GWs in Sec. 2.3, outlining the Post-Newtonian approximation for weak-field gravity in Subsec. 2.3.1, and numerical relativity for use in strong-field gravity in Subsec. 2.3.2. In Subsec. 2.4, we describe the effects of black hole spin and spin precession upon the gravitational waveform, before outlining the main concepts behind Bayesian inference with GW data for BBH parameters in Sec. 2.5. Finally, we summarise the contents of this chapter in Sec. 2.6. Throughout this thesis, we shall use Greek letter indices to represent four-vectors, such as x_μ , and Latin indices to represent 3-dimensional spatial vectors x^i . Except in Sec. 2.1 where we use SI units for G and c , in the remainder of this thesis we use units natural units $G = c = 1$ unless explicitly specified.

2.1 Introduction to Gravitational Waves in General Relativity

In this section, we first introduce the Einstein field equations, before going on to derive GWs in linearised gravity in Subsec. 2.1.1, and explicitly considering the source term in Subsec. 2.1.2. Throughout this section we closely follow the explanations presented in Refs. [41–43], which we refer the reader to for more details.

The Einstein equations of General Relativity (GR) are given as

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi GT_{\mu\nu}, \quad (2.1)$$

where $G_{\mu\nu}$ is the Einstein tensor, $R_{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar, G is the gravitational constant and $T_{\mu\nu}$ is the energy-momentum tensor. This is the key equation of Einstein's Theory of General Relativity, which is the best description we have of gravity. These equations may look elegant in the form above, but are deceptive in that they are actually 10 highly nonlinear coupled partial differential equations, with very few exact analytic solutions except in cases of high degrees of symmetry. In the ensuing subsections, we will show that GWs are a natural consequence of GR, and will derive the relevant expressions for them to motivate our study of merging binary black holes (BBHs) as astrophysical sources of these waves.

2.1.1 Gravitational Waves in Linearised Gravity

First we begin by considering GWs in linearised gravity, following the explanations of [41–43]. We consider a weak gravity approximation, in the form of a small deviation of the spacetime metric from the flat Minkowski metric,

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad (2.2)$$

where the metric perturbation $\|h_{\mu\nu}\| \ll 1$, ie. the magnitude of non-zero components of $h_{\mu\nu}$ is small, and $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$. We assume that since the perturbation is small, indices are raised and lowered with the flat metric $\eta_{\mu\nu}$ as usual, which is equivalent to assuming that the GWs do not carry enough energy and momentum to affect their own propagation. The expression for the inverse metric is

$$g^{\mu\nu} = \eta^{\mu\nu} - h^{\mu\nu}. \quad (2.3)$$

We can then construct the components of the Levi-Civita connection (or Christoffel symbols) which are

$$\Gamma_{\mu\nu}^{\lambda} = \frac{1}{2}\eta^{\lambda\rho}(\partial_{\nu}h_{\rho\mu} + \partial_{\mu}h_{\rho\nu} - \partial_{\rho}h_{\mu\nu}) = \frac{1}{2}(\partial_{\nu}h_{\mu}^{\lambda} + \partial_{\mu}h_{\nu}^{\lambda} - \partial^{\lambda}h_{\mu\nu}). \quad (2.4)$$

We construct the Riemann tensor, and subsequently the Ricci tensor and Ricci curvature scalar as follows:

$$R_{\mu\rho\nu}^{\lambda} = \partial_{\rho}\Gamma_{\mu\nu}^{\lambda} - \partial_{\nu}\Gamma_{\mu\rho}^{\lambda} = \frac{1}{2}(\partial_{\rho}\partial_{\mu}h_{\nu}^{\lambda} + \partial_{\nu}\partial^{\lambda}h_{\mu\rho} - \partial_{\rho}\partial^{\lambda}h_{\mu\nu} - \partial_{\nu}\partial_{\mu}h_{\rho}^{\lambda}), \quad (2.5)$$

$$R_{\mu\nu} = R_{\mu\lambda\nu}^{\lambda} = \frac{1}{2}(\partial_{\lambda}\partial_{\mu}h_{\nu}^{\lambda} + \partial_{\nu}\partial^{\lambda}h_{\mu\lambda} - \square h_{\mu\nu} - \partial_{\mu}\partial_{\nu}h), \quad (2.6)$$

$$R = R_{\mu}^{\mu} = \partial_{\nu}\partial^{\mu}h_{\mu}^{\nu} - \square h, \quad (2.7)$$

where $h = h_{\mu}^{\mu}$ is the trace of the metric perturbation tensor and $\square = \partial_{\mu}\partial^{\mu} = \eta^{\mu\nu}\partial_{\mu}\partial_{\nu}$ is the flat space d'Alembertian operator. We may now build the Einstein tensor,

$$G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}R = \frac{1}{2}(\partial_{\lambda}\partial_{\nu}h_{\mu}^{\lambda} + \partial^{\lambda}\partial_{\mu}h_{\nu\lambda} - \square h_{\mu\nu} - \partial_{\mu}\partial_{\nu}h - \eta_{\mu\nu}\partial_{\lambda}\partial^{\rho}h_{\rho}^{\lambda} + \eta_{\mu\nu}\square h), \quad (2.8)$$

This expression can be simplified using an appropriate gauge transform, which in General Relativity corresponds to a coordinate transform. The transformation law for tensors under coordinate transforms reads, for our metric tensor,

$$h'_{\mu\nu}(x'^{\sigma}) = \frac{\partial x^{\alpha}}{\partial x'^{\mu}} \frac{\partial x^{\beta}}{\partial x'^{\nu}} h_{\alpha\beta}(x^{\sigma}). \quad (2.9)$$

This transformation can be used with the infinitesimal coordinate transform

$$x^{\mu} \rightarrow x'^{\mu} = x^{\mu} + \zeta^{\mu}, \quad (2.10)$$

to give, to leading order in ζ ,

$$h_{\mu\nu} \rightarrow h'_{\mu\nu} = h_{\mu\nu} - \partial_\mu \zeta_\nu - \partial_\nu \zeta_\mu. \quad (2.11)$$

Therefore, if $h_{\mu\nu}$ is a solution of the linearised Einstein equations, then so too will $h'_{\mu\nu}$ be, for any $h'_{\mu\nu}$ given by (2.11). We may therefore now choose to work in a convenient coordinate system, known as harmonic gauge or de Donder gauge, which satisfies

$$g^{\mu\nu} \Gamma_{\mu\nu}^\lambda = 0. \quad (2.12)$$

This condition, in conjunction with the Christoffel symbol definition of Eq. (2.4) and the inverse metric definition of Eq. (2.3), gives

$$\partial_\mu h_\nu^\mu = \frac{1}{2} \partial_\nu h_\mu^\mu. \quad (2.13)$$

Using the expression of Eq. (2.8) and the condition of Eq. (2.13), the Einstein equations (2.1) then become,

$$\square h_{\mu\nu} - \frac{1}{2} \square h \eta_{\mu\nu} = -16\pi G T_{\mu\nu}. \quad (2.14)$$

This can be further simplified by transforming to the trace-reversed metric perturbation,

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} h, \quad (2.15)$$

so named because taking the trace of both sides gives $\bar{h} = \eta^{\mu\nu} \bar{h}_{\mu\nu} = \bar{h}_\mu^\mu = -h$. Therefore we can reconstruct the original metric tensor as

$$h_{\mu\nu} = \bar{h}_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} \bar{h}, \quad (2.16)$$

and the harmonic gauge condition of Eq. (2.13) in terms of the trace-reversed metric perturbation becomes

$$\partial_\mu \bar{h}^\mu_\nu = 0. \quad (2.17)$$

Rewriting Eq. (2.14) in terms of \bar{h} , we obtain

$$\square \bar{h}_{\mu\nu} = -16\pi G T_{\mu\nu}, \quad (2.18)$$

which may be recognised as a wave equation with a source term expressed in terms of the energy momentum tensor, $T_{\mu\nu}$. Gravitational waves are generated by this matter source term, however it is instructive to first consider the propagation of gravitational waves outside the source, in vacuum, where we have $T_{\mu\nu} = 0$. Then Eq. (2.18) becomes

$$\square \bar{h}_{\mu\nu} = 0. \quad (2.19)$$

We now exploit the remaining gauge freedom, and impose that $\bar{h} = 0$, which implies that $\bar{h}_{\mu\nu} = h_{\mu\nu}$, and we also impose that $h^{0i} = 0$. Therefore we have from the condition of Eq. (2.17) that

$$\partial^0 h_{00} = 0, \quad (2.20)$$

which tells us that the time-varying part of the metric perturbation, ie. the radiating GW, is contained within the spatial components of the metric h_{ij} . We are now in transverse-traceless (TT) gauge, which is defined by the equations

$$h^{0\mu} = 0, \quad (2.21)$$

$$h^i_i = 0, \quad (2.22)$$

$$\partial^i h^{ij} = 0, \quad (2.23)$$

where the first two equations are a result of the gauge we have imposed, and the third is as a result of the condition in Eq. (2.17). The general solutions to Eq. (2.19) in this

gauge are

$$h_{ij}^{\text{TT}} = e_{ij} e^{ik_\mu x^\mu}, \quad (2.24)$$

where $k^\mu = (\omega, k^i)$ is the wave vector, and e_{ij} is the polarisation tensor. We may choose that the wave propagates along the z -axis, which gives

$$h_{\mu\nu}^{\text{TT}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (2.25)$$

where h_+ and h_\times are the two independent GW polarisations. We can see from this that the GW perturbs the spacetime in a plane perpendicular to its direction of travel.

2.1.2 Sourcing the Gravitational Waves

This derivation of the GW signal in TT gauge assumed that our source term $T_{\mu\nu}$ was zero, however this doesn't give us any information as to how the GWs are actually generated by the source. Therefore, we will return to Eq. (2.18) but now assume a non-zero source term. As we are only considering weak-field gravity here, we assume that the components of $T_{\mu\nu}$ are small. We also assume that the source matter is localised to some spatial region Σ , and that we wish to calculate what the metric perturbation looks like far away from the region Σ . These approximations are reasonable for the case of an orbiting binary black hole source, for example. Outside of the region Σ , the energy-momentum tensor vanishes, and so we have $T_{\mu\nu}(t, \vec{x}') = 0$ for $\vec{x}' \notin \Sigma$. In the following, we compute the leading-order contribution to the spatial components of $h_{\mu\nu}$ for sources whose internal motions are slow compared to the speed of light. We follow closely the description in Ref. [44].

With the assumptions specified, we may write the solution for $\bar{h}_{\mu\nu}$ outside of Σ as usual for a wave equation using a Green's function $G(t, \vec{x}; t', \vec{x}')$, which describes the field

that arises due to a delta function source:

$$\square G(t, \vec{x}; t', \vec{x}') = \delta(t - t')\delta(\vec{x} - \vec{x}'). \quad (2.26)$$

We can then express the field arising from our actual source by substituting into Eq. (2.18),

$$\bar{h}_{\mu\nu}(t, \vec{x}) = -16\pi \int dt' d\vec{x}' G(t, \vec{x}; t', \vec{x}') T_{\mu\nu}(t', \vec{x}'), \quad (2.27)$$

where the form of the Green's function associated with the operator \square is known to be

$$G(t, \vec{x}; t', \vec{x}') = -\frac{\delta(t' - [t - |\vec{x} - \vec{x}'|])}{4\pi|\vec{x} - \vec{x}'|}. \quad (2.28)$$

The quantity given by $[t - |\vec{x} - \vec{x}'|]$ is the retarded time, which describes the lag between propagation of information from events at \vec{x}' to an observer at \vec{x} . We therefore have

$$\bar{h}_{\mu\nu}(t, \vec{x}) = 4 \int d\vec{x}' \frac{T_{\mu\nu}(t - |\vec{x} - \vec{x}'|, \vec{x}')}{|\vec{x} - \vec{x}'|}. \quad (2.29)$$

We wish to evaluate this equation at large distances from the source, where we may Taylor expand

$$\frac{1}{|\vec{x} - \vec{x}'|} \approx \frac{1}{r} + \frac{x^i x'^i}{r^3} + \dots, \quad (2.30)$$

where we have introduced $r = |\vec{x}|$. The spatial portion of Eq. (2.29) then becomes

$$\bar{h}_{ij}(\vec{x}) = \frac{4}{r} \int d\vec{x}' T_{ij}(t - r, \vec{x}'). \quad (2.31)$$

In linearised theory, the energy-momentum tensor will adhere to the flat-space conservation law

$$\partial_\mu T^{\mu\nu} = 0, \quad (2.32)$$

which means that sources will move on geodesics of flat Minkowski space. Splitting this

up into time and spatial components we have

$$\partial_t T^{tt} + \partial_i T^{ti} = 0, \quad (2.33)$$

$$\partial_t T^{ti} + \partial_j T^{ij} = 0, \quad (2.34)$$

and so it follows that

$$\partial_t^2 T^{tt} = \partial_k \partial_l T^{kl}. \quad (2.35)$$

We may multiply both sides of the equation by $x^i x^j$, and manipulate the right hand side to obtain

$$\partial_t^2 (T^{tt} x^i x^j) = \partial_k \partial_l (T^{kl} x^i x^j) - 2\partial_k (T^{ik} x^j + T^{kj} x^i) + 2T^{ij}. \quad (2.36)$$

Inserting this into Eq. (2.31) then gives

$$\bar{h}_{ij} = \frac{4}{r} \int d\vec{x}' \left[\frac{1}{2} \partial_t^2 (T^{tt} x'^i x'^j) + \partial_k (T^{ik} x'^j + T^{kj} x'^i) - \frac{1}{2} \partial_k \partial_l (T^{kl} x'^i x'^j) \right] \quad (2.37)$$

$$= \frac{2}{r} \frac{\partial^2}{\partial t^2} \int d\vec{x}' \rho x'^i x'^j, \quad (2.38)$$

where we have used Gauss' theorem to recast the second and third terms in the integrand as surface integrals, which, taking the surface to be outside the source, reduce to zero. We have also used the fact that T^{tt} is just the mass density ρ . We may then define the moment of inertia tensor I_{ij} to be

$$I_{ij}(t) = \int d\vec{x}' \rho(t, \vec{x}') x'^i x'^j, \quad (2.39)$$

and, inserting this into Eq. (2.38), obtain

$$\bar{h}_{ij} = \frac{2}{r} \frac{\partial^2}{\partial t^2} (I_{ij}(t - r)). \quad (2.40)$$

This result describes GWs generated from a weak source, as observed far away from that source. It shows that the GWs are produced by the acceleration of masses, and that the

GW radiation is quadrupolar in nature, as no monopole or dipole solutions are admitted by the conservation law for the energy momentum tensor in Eq. (2.32). Now that we have derived the required expressions for GWs within linearised gravity, we will now consider that astrophysical sources which produce them in the next section.

2.2 Astrophysical Sources of Gravitational Waves

In the following subsections, we introduce the most important astrophysical sources of gravitational waves. We first introduce compact binary coalescences in Subsec. 2.2.1, which are the main focus of this thesis. We then for completeness introduce other promising astrophysical sources of GWs in Subsec. 2.2.3.

2.2.1 Compact Binary Coalescences

The focus of this thesis, and the only sources from which GWs have been detected so far [45, 134, 218, 255], are the mergers of compact binaries. These binaries are made up of compact objects (black holes, neutron stars or white dwarfs) in some combination, and their dynamics is dominated by gravitation. We note that the quadrupole formula of Eq. (2.40) was derived assuming a flat background, as the conservation law of Eq. (2.32) forces bodies to move on geodesics of flat Minkowski spacetimes. This assumption is not applicable to systems whose dynamics are governed by self gravity such as compact binary systems, and there will be a back reaction on the background spacetime due to the binary motion. It can be shown, however, that the expression still holds even for such self-gravitating systems like compact binaries, provided the binary is widely-separated. For further explanation we refer the reader to Ref. [44], for example. To lowest order this approximation will describe a binary of point particles with Newtonian dynamics in a flat Minkowski spacetime background. We assume the binary is moving on a circular orbit in the xy -plane with orbital separation R and orbital frequency Ω , as shown in Fig. 2.1. The orbital angular momentum \vec{L} of the binary is parallel to the z -axes in this coordinate




Figure 2.1: A BBH moving on a circular orbit. The orbital separation of the two black holes, denoted by the black dotted line, is R , and the masses of the two black holes are m_1 and m_2 , where we assume that $m_1 \geq m_2$. The edge of the purple circle indicates the path of the binary orbit, which, as the BHs are assumed to be non-spinning here, is confined to the xy -plane.

system.

If we assume that the binary is of equal mass $m_1 = m_2$, then the motion of each object is given by

$$\vec{x}_1(t) = \begin{pmatrix} R \cos(\Omega t)/2 \\ R \sin(\Omega t)/2 \\ 0 \end{pmatrix}, \quad (2.41)$$

$$\vec{x}_2(t) = \begin{pmatrix} -R \cos(\Omega t)/2 \\ -R \sin(\Omega t)/2 \\ 0 \end{pmatrix} \quad (2.42)$$

In the centre-of-mass frame of the binary the mass density ρ for our point-particle

binary is

$$\rho(\vec{x}, t) = M\delta(z)\left[\delta\left(x - \frac{R\cos(\Omega t)}{2}\right)\delta\left(y - \frac{R\sin(\Omega t)}{2}\right) + \delta\left(x + \frac{R\cos(\Omega t)}{2}\right)\delta\left(y + \frac{R\sin(\Omega t)}{2}\right)\right], \quad (2.43)$$

and so we may compute the components of I_{ij} according to Eq. (2.39),

$$I_{ij} = \begin{pmatrix} R^2 \cos^2(\Omega t) & R^2 \sin(\Omega t) \cos(\Omega t) & 0 \\ R^2 \sin(\Omega t) \cos(\Omega t) & R^2 \sin^2(\Omega t) & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (2.44)$$

$$= \frac{R^2}{2} \begin{pmatrix} (1 + \cos(2\Omega t)) & \sin(2\Omega t) & 0 \\ \sin(2\Omega t) & (1 - \cos(2\Omega t)) & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (2.45)$$

which, using Eq. (2.40) gives

$$\bar{h}_{ij} = -\frac{4MR^2\Omega^2}{r} \begin{pmatrix} \cos(2\Omega(t-r)) & \sin(2\Omega(t-r)) & 0 \\ \sin(2\Omega(t-r)) & -\cos(2\Omega(t-r)) & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (2.46)$$

where r is the distance to the binary from the observer. If we now consider a binary where the two objects are not equal mass, this generalises to

$$\bar{h}_{ij} = -\frac{4M\eta R^2\Omega^2}{r} \begin{pmatrix} \cos(2\Omega(t-r)) & \sin(2\Omega(t-r)) & 0 \\ \sin(2\Omega(t-r)) & -\cos(2\Omega(t-r)) & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (2.47)$$

where η is the symmetric mass ratio

$$\eta = \frac{\mu}{M}, \eta \leq \frac{1}{4}, \quad (2.48)$$

μ is the binary reduced mass

$$\mu = \frac{m_1 m_2}{M}, \quad (2.49)$$

and M is the total mass

$$M = m_1 + m_2. \quad (2.50)$$

We also introduce here the mass ratio q ,

$$q = \frac{m_1}{m_2} \geq 1. \quad (2.51)$$

In the Newtonian limit, the energy of a circular orbit is given by

$$E = -\frac{\eta M^2}{2R}, \quad (2.52)$$

and the orbital frequency Ω is given by

$$\Omega^2 = \frac{2GM}{R^3}. \quad (2.53)$$

Therefore, we may express the gravitational wave polarisations in Eq. (2.25) as

$$h_+(t-r) = -\frac{2G\mu M}{rR} \cos(2\Omega(t-r)), \quad (2.54)$$

$$h_\times(t-r) = -\frac{2G\mu M}{rR} \sin(2\Omega(t-r)). \quad (2.55)$$

The GW flux is given by the quadrupolar formula

$$\mathcal{F} = \frac{G}{5c^5} \langle \ddot{I}_{ij} \ddot{I}_{ij} \rangle \quad (2.56)$$

where $\ddot{I}_{ij} = \frac{\partial^3}{\partial t^3} I_{ij}$, and so using our previously derived expression for the moment of inertia tensor in Eq. (2.45), we have

$$\mathcal{F} = \frac{32\eta^2 M^2}{5R^5} = \frac{dE}{dt} \quad (2.57)$$

where, explicitly, we have set $G = c = 1$. We can therefore calculate the inspiral rate

$$\frac{dR}{dt} = \frac{dR}{dE} \frac{dE}{dt} = -\frac{64}{5} \frac{\mu M^2}{R^3}. \quad (2.58)$$

As the emitted GWs are quadrupolar in nature, they will have frequency $f_{GW} = \Omega/\pi$. Combining Eqs. (2.53) and (2.58) with this new definition for gravitational wave frequency, we obtain

$$\frac{df_{GW}}{dt} = \frac{96}{5} \pi^{8/3} \mathcal{M}^{5/3} f_{GW}^{11/3}, \quad (2.59)$$

where \mathcal{M} is the chirp mass, given by

$$\mathcal{M} = \mu^{3/5} M^{2/5} \quad (2.60)$$

Eq. (2.59) is the characteristic frequency evolution for a CBC inspiral known as a chirp signal.

In light of this derivation of the frequency evolution for a compact binary, we broadly summarise the phenomenology of binary black hole mergers in three stages: inspiral, merger, and ringdown. In the inspiral, as the binary evolves, it loses binding energy due to the emission of GWs, so the orbital separation shrinks. The frequency of the orbit, and resulting GW frequency, increases according to Eq. (2.59), and the amplitude of the emitted GW signal also increases as the binary inspirals, as can be seen from Eq. (2.56) with decreasing separation R . This process continues until the binary eventually plunges and merges into a single object, at which point the GW signal is at its maximum amplitude. This merger could directly form a black hole, or if the binary masses are light enough, they may form a stable neutron star, or hypermassive neutron star which then collapses into a black hole [46]. If the remnant object is a black hole, immediately after merger it will behave as a perturbed single black hole, and will settle down to a stable individual black hole configuration. The remnant BH in this ringdown stage will emit GWs in a spectrum of quasi-normal modes [47], whose frequencies depend only on




Figure 2.2: Polarizations h_+ , h_\times of a BBH coalescence showing the inspiral, merger and ringdown, aligned such that the merger occurs at a time of 0 seconds. The BBH has parameters $\{M, q, \chi_1, \chi_2, d_L, \iota, \psi, \Phi_0\} = \{100M_\odot, 1, \vec{0}, \vec{0}, 100\text{Mpc}, 0, 0, 0\}$ (see Sec. 2.2.2 for parameter definitions), where Φ_0 is quoted at a reference frequency of 4Hz, and shows the monotonically increasing amplitude and phase in the inspiral. At merger the amplitude reaches maximum, before exponentially decaying in the ringdown where the remnant black hole is settling down. The waveforms were created using the SEOBNRv4PHM [48] waveform model within LALSimulation [49].

the mass and spin angular momentum of the BH. A full inspiral-merger-ringdown (IMR) signal of a BBH merger is depicted in Fig. 2.2, where the waveform polarizations $h_+(t)$ and $h_\times(t)$ are aligned such that the merger occurs at a time of 0 seconds. The inspiral before the merger at $t = 0$ clearly shows the behaviour of monotonically increasing amplitude and phase. At $t = 0$, the merger occurs and the waveform amplitude reaches its peak. In the ringdown ($t > 0$) the amplitude decays exponentially as the remnant BH is settling down and shedding excess energy and angular momentum through GWs as a spectrum of quasinormal modes [47].

2.2.2 Detecting Gravitational Waves from Black Hole Mergers

In Subsec. 2.2.1, we derived the expression in Eq. (2.47) for a GW signal of an inspiralling compact binary in the Newtonian limit. However, we did not take into account the response of a GW detector, or the distance or orientation of the source with respect to us as observers. In this section we consider the sensitivity of the detector, and how the parameters of the binary affect the signal we observe, including those parameters intrinsic to the binary itself, and those which describe its position and orientations relative to an observer.

GW interferometers detect signals using extremely small phase changes between two interfering lasers, where the fractional change for a typical CBC signal is of the order 10^{-21} . In order to effectively detect GW signals, the noise sources in a detector which can cause phase changes must be well understood, so as not to confuse noise with an astrophysical signal. Sources of noise for ground-based detectors include seismic noise, thermal noise, shot noise, radiation pressure noise, gravity gradient noise, and many other sources. We refer the reader to Ref. [50] for a more thorough discussion of noise sources and extraction of GW signals from detector data. The sum of these noise source contributions defines the detector sensitivity, characterised by the *power spectral density* (PSD) (see Fig. 6.1 for PSDs of the planned Advanced LIGO A+ upgrade [51], and future ground-based instruments Cosmic Explorer [52] and the Einstein Telescope [53]).

Consider an L-shaped GW detector in a 2-dimensional coordinate system where the axes are defined as being parallel to the detector arms, as shown in Fig. 2.3. The detector lies in the xy -plane, with arms positioned at the location of the x and y -axis as shown in the Figure. The location of the binary on the sky, as seen from the detector, is denoted by the two angles θ and ϕ , and the polarisation angle ψ denotes the angle between the direction of the $+$ -polarisation and the observer's line of sight, projected in the horizontal plane of the detector.

We may express the gravitational wave polarisations $h_+(t)$ and $h_\times(t)$ in the Newtonian

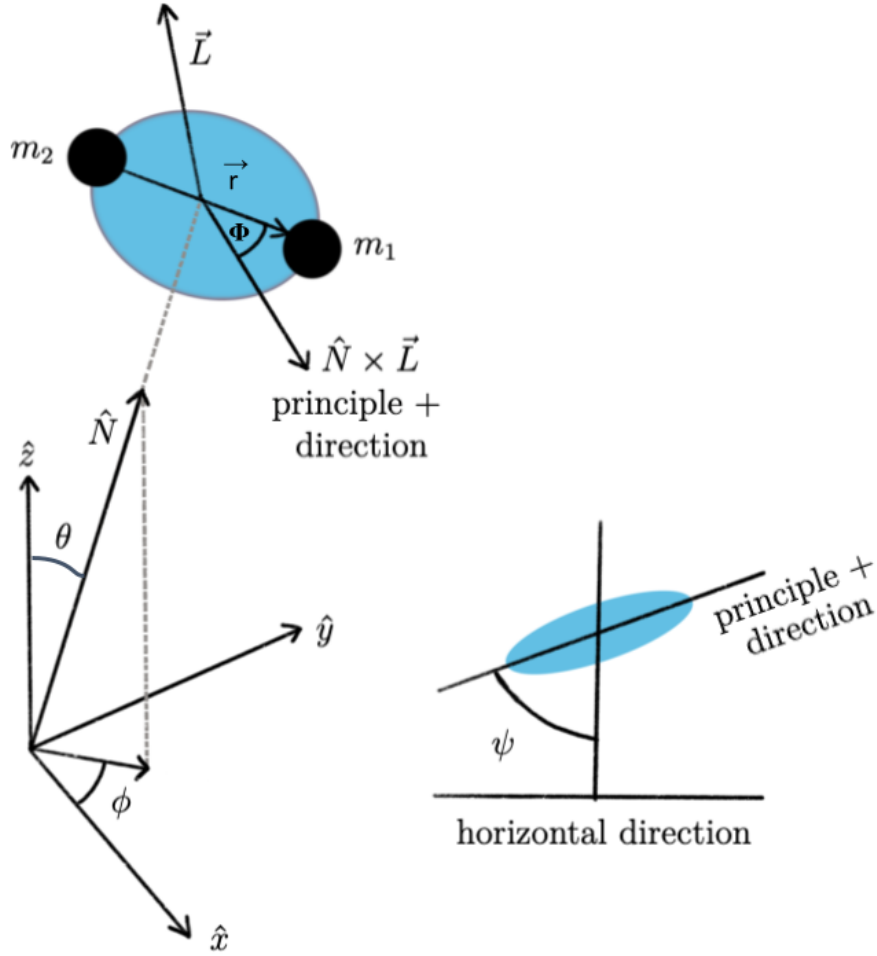


Figure 2.3: BBH with relative orientation to observer. The L-shaped GW detector lies within the xy -plane, and the location and orientation of the binary with respect to this detector are characterised by the angles θ and ϕ to represent the sky location, ψ the polarisation, and Φ the binary orbital phase. Fig. adapted from Ref. [14].

limit as

$$h_+(t) = -\frac{2\mu M}{rD} \left[1 + (\hat{L} \cdot \hat{N})^2 \right] \cos(2\Phi t), \quad (2.61)$$

$$h_\times(t) = -\frac{2\mu M}{rD} \left[-2\hat{L} \cdot \hat{N} \right] \sin(2\Phi t), \quad (2.62)$$

where we note that the only difference between these expressions and those of Eqs. (2.54) are the factors in square brackets which represent the misalignment of the plane of the binary orbit with respect to the line of sight of the observer. We parameterise this using

the binary inclination ι , which is defined by

$$\cos(\iota) = \hat{L} \cdot \hat{N}. \quad (2.63)$$

The detector strain $h(t)$ that the GWs produce in the interferometer can be described as the linear combination of the two polarisations $h_+(t)$, $h_\times(t)$ as follows,

$$h(t) = F_+(\theta, \phi, \psi)h_+(t) + F_\times(\theta, \phi, \psi)h_\times(t), \quad (2.64)$$

where F_+ , F_\times are the detector beam pattern coefficients. These are geometric functions that describe the strength of detector response to a GW signal, depending on the location of the source in the sky, and its relative orientation to the plane of the detector, denoted by the angles θ , ϕ and ψ in Fig. 2.3.

In Eq. (2.59) we derived the frequency evolution for a non-spinning, inspiralling BBH. We note that both Eq. (2.59) and also Eqs. (2.61), (2.62) are invariant under the transformation

$$(f, \mathcal{M}, \mu, r, R, t) \Rightarrow (f/\lambda, \mathcal{M}\lambda, \mu\lambda, r\lambda, R\lambda, t\lambda). \quad (2.65)$$

Therefore we can derive from the GW signal the redshifted masses represented by $\mathcal{M}\lambda$, often called detector-frame masses, where

$$\mathcal{M}_{\text{detector}} = (1 + z)\mathcal{M}_{\text{source}}, \quad (2.66)$$

where $\mathcal{M}_{\text{source}}$ is the true chirp mass as measured in the frame of the binary source, and z is the redshift. Similar relations will hold for the other binary mass parameters. This redshift value z will have an associated luminosity distance d_L (see Eq. (6.36)). We note that as we directly measure the detector frame chirp mass $\mathcal{M}_{\text{detector}}$, which will be larger than $\mathcal{M}_{\text{source}}$, the relation in Eq. (2.65) implies that the frequency of the GW signal will be redshifted to appear lower. This mass-frequency scaling is a general feature of GWs from CBCs, that the merger of binaries with higher total masses occurs at lower frequencies

than binaries with smaller masses. Consequently when building gravitational waveform models as discussed further in Chap. 3, we can treat the total mass M as an overall scaling factor.

Therefore a GW signal from a merging BBH depends upon 16 parameters:

$$h(t; \vec{\theta}) = h(t; m_1, m_2, \vec{\chi}_1, \vec{\chi}_2, e, d_L, \alpha, \delta, \iota, \psi, t_c, \phi_c), \quad (2.67)$$

where m_1 and m_2 are the masses of the two BHs with $m_1 \geq m_2$, which may be alternatively parameterised by any pair combination of: the total mass M from Eq. (2.50); the chirp mass \mathcal{M} of Eq. (2.60); the reduced mass μ from Eq. (2.49); the symmetric mass ratio η from Eq. (2.48); or the mass ratio q given from Eq. (2.51).

The vectors $\vec{\chi}_1$ and $\vec{\chi}_2$ are the normalised 3-dimensional spin vectors of the individual black holes, defined as

$$\vec{\chi}_i = \frac{\vec{S}_i}{m_i^2} < 1, \quad (2.68)$$

where the inequality is as a result of the Kerr limit for BH spins. The orbital eccentricity e we will hereafter assume to be zero. This is because, while the binary may have formed on an eccentric orbit, the radiation of GWs throughout the early inspiral efficiently removes energy and angular momentum from the system, and so most binaries are expected to circularise by the time they reach the frequency band of ground-based detectors [54]. These quantities make up the intrinsic parameters of the binary, as they are properties of the binary itself rather than its position and orientation with respect to an observer. For convenience, we define the set of 7 intrinsic parameters required for modelling a binary with generically-oriented spins as

$$\vec{\theta} = \{q, \vec{\chi}_1, \vec{\chi}_2\}, \quad (2.69)$$

where we have not included the total mass M as it acts only as a scaling factor upon the waveform.

The remaining parameters are the extrinsic parameters, where d_L is the luminosity distance between the binary and the observer, ψ is the polarisation angle from Fig. 2.3, and α and δ are the right ascension and declination respectively which represent the location of the binary source on the sky from the observer's perspective. The inclination ι is the angle between the binary orbital angular momentum \vec{L} and the observer's line of sight \vec{N} as seen in Fig. 2.3, defined by

$$\cos(\iota) = \hat{L} \cdot \hat{N}. \quad (2.70)$$

This orientation dependence can be alternatively parameterised by θ_{JN} . This is defined as the angle between \hat{N} and the binary total angular momentum \vec{J} , which is the sum of the orbital angular and the individual spin angular momenta,

$$\vec{J} = \vec{L} + \vec{S}_1 + \vec{S}_2, \quad (2.71)$$

and we have that

$$\cos(\theta_{JN}) = \hat{J} \cdot \hat{N}. \quad (2.72)$$

Finally, t_c and ϕ_c are the time and binary phase at the time of coalescence. We note that each of these parameters will imprint upon the GW signal in different ways, and there exists degeneracies between the effects of different parameters upon the waveform. One example is the degeneracy between mass ratio q and the components of BH spin aligned with \vec{L} , which we discuss in more detail in Subsec. 2.2.2, and a further example is provided in Sec. 3.1 between d_L and ι .

Throughout this thesis, we will switch between descriptions of the gravitational waveform in the time domain and frequency (Fourier) domain, depending on convenience. Therefore we define here the Fourier transform for a function $f(t)$ as

$$\mathcal{F}(f(t)) = \int_{-\infty}^{\infty} f(t) e^{2\pi i f t} dt = \tilde{f}(t). \quad (2.73)$$

We will also often require a method to quantify the agreement between two waveforms, so we will employ the (frequency-domain, noise-weighted) inner product or match \mathcal{M}_f optimised over a time and phase shift, given as

$$\mathcal{M}_f(h_1, h_2) = \max_{t_c, \phi_0} \frac{\langle h_1, h_2 \rangle}{\sqrt{\langle h_1, h_1 \rangle \langle h_2, h_2 \rangle}}, \quad (2.74)$$

where the inner product is defined as

$$\langle h_1, h_2 \rangle = 4\Re \int_{f_{\min}}^{f_{\max}} \frac{\tilde{h}_1(f) \tilde{h}_2^*(f)}{S_n(|f|)} df, \quad (2.75)$$

with $S_n(f)$ the one-sided PSD of the detector noise, \tilde{h} indicates the Fourier transform of h , and \tilde{h}^* the complex conjugate. The mismatch can now be defined as

$$\bar{\mathcal{M}}_f(h_1, h_2) \equiv 1 - \mathcal{M}_f(h_1, h_2). \quad (2.76)$$

We will also find it convenient to introduce a normalized waveform $\hat{h} = h / \sqrt{\langle h, h \rangle}$.

When using a white-noise PSD, i.e. independent of frequency, it is convenient to define a time-domain overlap

$$\langle h_1, h_2 \rangle_t = \Re \int_{t_{\min}}^{t_{\max}} h_1(t) h_2^*(t) dt, \quad (2.77)$$

with an inherited norm $\|h\|_t^2 = \langle h, h \rangle_t$. We can then define an analogous time domain match as

$$\mathcal{M}_t(h_1, h_2) = \frac{\langle h_1, h_2 \rangle_t}{\|h_1\|_t \|h_2\|_t}, \quad (2.78)$$

and the associated *mismatch* as $\bar{\mathcal{M}}_t = 1 - \mathcal{M}_t$.

2.2.3 Other Astrophysical Sources of Gravitational Waves

Aside from the mergers of compact binaries, there are several other astrophysical phenomena which are expected to produce GW signals observable by current or future ground-based detectors. We highlight a few key examples here for completeness, noting their morphology and how it differs from CBC signals, but for a more thorough review we refer the reader to, for example, Refs. [41, 55].

When a massive star with mass $M_* \gtrsim 10M_\odot$ comes to the end of its life, it can no longer support itself with outward radiation pressure, and so undergoes gravitational collapse. Matter from the outer layers of the dying star descend radially inwards onto the core, after which it is expelled rapidly in a supernova explosion. This satisfies many of the requirements for a strong GW emission, as this large amount of mass in a relatively concentrated volume of space reaches relativistic speeds. The degree of asymmetry in the explosion will also determine the strength of the emitted signal, though mechanisms for producing non-spherical emission may include rapid rotation of the star's core leading to instabilities during collapse, or the presence of magnetic fields within the star. The resulting GW signal will be a short burst, which may be modelled by a sum of sine Gaussians.

Another key theorised source of GW signals is from rapidly spinning neutron stars, known as pulsars, with some degree of asymmetry in their mass distribution. This asymmetry will induce the time-varying quadrupole moment and lead to continuous GW emission at a single frequency, so the pulsar will lose energy and angular momentum (it will 'spin down'). Non-axisymmetry could be caused by oblateness misaligned with the pulsar's axis of rotation, for example due to misalignment of the star's internal magnetic field with this axis, or as a result of accreting material from a companion star. As in the core collapse case above, the strength of the GW emission will be proportional to the degree of asymmetry in the source. Since no GW signals from known pulsars have been detected thus far, this allows us to place upper limits on the asymmetry in the mass distribution of these sources, for example in Ref. [56].

Finally, in addition to individually resolvable sources of GW signals, we also expect a stochastic background of GWs which may span across a wide range of frequencies. This background is thought to contain contributions from a variety of sources, including early universe components analogous to the cosmic microwave background of photons. Phenomena such as early universe phase transitions, signatures of cosmic strings, and amplification of primordial density fluctuations through inflation are all expected to produce GW background contributions, though no detection of a stochastic background has yet been made.

2.3 Modelling Gravitational Waveforms

As previously mentioned in Sec. 2.2.1, the typical gravitational waveform from a CBC can be divided into three main regions: the inspiral, merger, and ringdown, as shown in Fig. 2.2. In the inspiral, where the separation of the binary is large and the orbital velocity is small, we need not solve the full nonlinear Einstein field equations exactly, and instead use a post-Newtonian (PN) expansion [57] method to derive accurate expressions for the waveform. In the strong-field (later inspiral, plunge, merger and ringdown), the PN approximation breaks down as the orbital separation shrinks and velocities become highly relativistic, and we must use numerical methods to computationally solve the full Einstein field equations with numerical relativity. In the following subsections, we briefly summarise each of these methods. We note that these approximations are only the building blocks of state-of-the-art waveform approximants, in Chapter 3 we describe in more detail how these pieces are used to create accurate waveforms of a full IMR signal.

2.3.1 Weak Field Solutions and the Post-Newtonian Approximation

Generation of GWs in linearised gravity (as seen in Sec. 2.1) assumes that the background spacetime can be taken as flat, thereby assuming that the system producing GWs con-

tributes negligibly to the spacetime curvature near to the source, and so the background spacetime curvature and source velocity can be treated as independent. This is a reasonable approximation for systems whose dynamics is not governed by gravitational forces, but this is not the case for CBCs, as these systems are held together by gravity (self-gravitating) and so the spacetime curvature and source velocities are not independent: the binary motion is affected by the spacetime curvature produced by its own motion. In this case, we need to take into account the curvature of the background spacetime near the source due to its motion.

We consider a compact binary source which is slowly-moving and weakly self-gravitating, such that $v/c \ll 1$ where v is the characteristic source velocity (in this the orbital velocity of the black holes), and $R_s/d \ll 1$ where R_S is the Schwarzschild radius and d is the typical size of the system (the black hole separation).

In the early inspiral, the two bodies are far apart such that their separation R is much larger than the intrinsic characteristic scale of the binary set by the total mass M . In this regime, the binary undergoes quasicircular motion of gradually decaying orbits due to the emission of GWs, and the system can be described as slowly-moving and weakly self-gravitating. Therefore, the motion can be approximately solved using a PN expansion in the characteristic velocity of the binary, $\frac{v}{c}$. Here we present only the fundamental ideas which govern this formalism, and for a more thorough treatment we refer the reader to the review article in Ref. [57].

We may schematically write the PN-expanded equations of motion for the binary as [15],

$$\frac{d^2 x^i}{dt^2} = -\frac{m x^i}{r^3} \left[1 + \mathcal{O}(\epsilon) + \mathcal{O}(\epsilon^{3/2}) + \mathcal{O}(\epsilon^2) + \mathcal{O}(\epsilon^{5/2}) + \dots \right] \quad (2.79)$$

where we introduce the parameter $\epsilon \equiv \sqrt{v/c}$ and have expanded up to (post)^{5/2}-Newtonian (2.5PN) order, where the maximum power of ϵ represents the order of expansion. The leading order term is the Newtonian limit, which we obtain by combining the expressions for the Newtonian polarisations in Eq. (2.47) and Eqs. (2.61), (2.61) into the complex

strain,

$$h(t) = h_+(t) - ih_\times(t), \quad (2.80)$$

to obtain

$$h(t) = -\mathcal{A} \frac{4M\nu v^2}{r} e^{-i(2\Omega t + \Phi_0)}, \quad (2.81)$$

where here \mathcal{A} is an amplitude factor which groups together the effects of the sky location, detector geometry and binary inclination as previously discussed in Sec. 2.2.2. Therefore the evolution of the binary and GW emission is governed by the orbital phase evolution, since

$$\Omega = \frac{d\Phi_{\text{orb}}}{dt} = \frac{v^3(t)}{M}, \quad (2.82)$$

and so the time evolution of the binary phase must also be predicted as an expansion in ϵ .

In the PN approximation, we assume that the emitted gravitational binding energy $E(v)$ and GW flux $\mathcal{F}(v)$ are expansions in the orbital velocity v/c . We have already seen the Newtonian limit for the energy of a circular orbit in Eq. (2.52), and the GW flux in Eq. (2.56). We assume that the binary undergoes circular orbits, which are gradually decaying, and so the velocity can be expressed using Kepler's law as

$$v = \sqrt{\frac{M}{R}} = (M\Omega)^{1/3}, \quad (2.83)$$

and so the energy and flux become

$$E_0(v) = -\frac{M\eta}{2}v^2, \quad (2.84)$$

$$\mathcal{F}_0(v) = -\frac{32}{5}\eta^2v^{10}, \quad (2.85)$$

where here the subscript 0 denotes that we are assuming Newtonian gravity for the background spacetime. More specifically, we assume in this lowest order approximation that the binary motion is not affected by the spacetime curvature produced by its own motion.

We may then define the expansions in v for energy and flux,

$$E(v, \vec{\theta}) = E_0 \left[1 + \sum_{i=2}^n E_i(\vec{\theta}) v^i + \ln(v) \sum_{i=8}^n E_i^{(l)}(\vec{\theta}) v^i \right] + \dots, \quad (2.86)$$

$$\mathcal{F}(v, \vec{\theta}) = \mathcal{F}_0 \left[1 + \sum_{i=2}^n \mathcal{F}_i(\vec{\theta}) v^i + \ln(v) \sum_{i=6}^n \mathcal{F}_i^{(l)}(\vec{\theta}) v^i \right] + \dots \quad (2.87)$$

where we are explicitly assuming units of $c = 1$, and $\vec{\theta}$ are the parameters of the binary. The quantities E_i and \mathcal{L}_i are the $i/2$ -order PN term coefficients in the PN expansions for flux and energy respectively, and we also have logarithmic terms in the expansions beginning at 4-PN order in the energy, and 3-PN order in the flux.

We make the assumption that the orbital decay is entirely due to GW emission, and therefore due to conservation of energy we have

$$\mathcal{F}(v) = -\frac{dE(v)}{dt}. \quad (2.88)$$

Furthermore, we assume that the binary undergoes an adiabatic inspiral, in which the fractional change of the orbital frequency over a single orbital period is small, and so we have $\dot{\Omega}/\Omega^2 \ll 1$. Therefore we can specify the phase evolution with the two equations:

$$\frac{d\Phi_{\text{orb}}(t)}{dt} = \frac{v^3}{M}, \quad (2.89)$$

$$\frac{dv(t)}{dt} = -\frac{\mathcal{F}(v)}{E'(v)}, \quad (2.90)$$

where $E'(v)$ denotes the derivative of E with respect to v . These differential equations can then be solved to obtain an expression for the waveform. The different approaches to solving these systems of equations are commonly referred to as PN approximants, for which a thorough discussion and comparison can be found in Ref. [58]. In order to define the energy and flux associated with the GWs themselves, a stress-energy tensor for the GWs must then be constructed. For an example of construction of a pseudo-tensor from second-order terms in $h_{\mu\nu}$, where the GWs are presumed to be perturbations of Minkowski

space, we refer the reader to Ref. [41].

Therefore, having seen in Sec. 2.1.1 that GWs are produced from a time varying quadrupole moment at leading order, the PN approximation now provides a description of the GWs in terms of higher order radiative multipoles [57, 59, 60]. These multipoles are coupled to the GW energy and flux, and are summed to produce the GWs in terms of h_{ij}^{TT} . The polarizations waveforms $h_{+,\times}$ are then computed using projections, and schematically we have

$$h_{+,\times}^{(n)} = \frac{2M\eta}{r} v^2 \sum_{n=0}^{\infty} v^n H_{+,\times}^{(n/2)}, \quad (2.91)$$

where n is the PN order, and the PN coefficients $H_{+,\times}$, which depend on the binary orbital phase Φ , are given for example in Ref. [57] up to 3.5PN for a non-spinning binary.

2.3.2 Numerical Relativity and the Strong Field Regime

In the previous subsection, we discussed the PN approximation, which can be used to compute waveforms in the inspiral phase of a BH merger, where the binary is slow-moving. However, as the binary separation decreases and the orbital velocity becomes more relativistic, the PN approximation breaks down. Therefore we must use the full nonlinear Einstein field equations, solving them numerically with numerical relativity (NR). In this subsection, we briefly discuss the main concepts behind NR as relevant to this thesis, and for a more thorough discussion we refer the reader to Refs. [61, 62], for example.

Numerical relativity is a computational approach to solve the full, nonlinear Einstein field equations, which formally depend on all ten degrees of freedom in the metric tensor $g_{\mu\nu}$. We are no longer assuming a weak perturbation $h_{\mu\nu}$ on a Minkowski metric as we are in the strong-field regime of gravity, and so this approach requires solving ten coupled, nonlinear, and partial differential equations. Introduction of a multidimensional grid coordinate system (usually one time-like and three space-like coordinates, referred to as 3 + 1 dimensions) allows for formulation of these ten equations into six dynamical

equations, and four constraint equations, commonly referred to as ADM equations after Arnowitt, Deser and Misner [63]. The dynamical equations describe the time evolution of the spatial portion of the metric tensor g_{ij} , and its time derivative \dot{g}_{ij} , given a set of initial conditions (often called initial data) which must satisfy the constraint equations. These $3 + 1$ equations are not unique, and one must choose a formulation which allows for stable evolution of the Einstein equations through time integration. An example of a reformulation which satisfies well-posedness of the equations for CBCs is the BSSN (Baumgarte, Shapiro, Shibata, Nakamura) [64, 65] formulation. It was introduced in the 1990s and introduces auxiliary variables that help to regularize the equations near singularities, allowing for more stable numerical solutions. We note that not all successful methods use a $3 + 1$ formulation, for example the generalised harmonic decomposition was used to complete the first successful simulation of a BBH merger, and extraction of the resulting GW signal, in 2005 by Pretorius [66]. Having formulated the equations, gauge freedom must then be fixed by specifying four variables: the lapse function and shift vector. The lapse function describes the proper time of an observer's normal to their current spatial hypersurface, and the 3-dimensional shift vector is the deviation of the observer's timeline from the normal to the current hypersurface, projected onto that hypersurface.

Beginning from a numerical solution to the constraint equations, the system is then evolved forwards in time until the merger using numerical integration methods. To deal with the curvature singularity of the BHs, approaches can be used such as BH excision [67], where the interior of the apparent horizon is removed from the computational domain, or a moving puncture method where slices are placed such that they do not include the BH singularity. Extraction of the GW signal should be performed as far away from the source as possible, which is obfuscated by a finite sized computational grid, so extrapolation methods of extraction at successively larger radii may be used. Finally, simulations of merging BBHs need to adequately cover both the areas of large curvature near the BH horizons, as well as far from the source where GW extraction takes place. Therefore

covering the whole domain with a single grid spacing is computationally infeasible, and so simulations often use multiple overlapping grids with fixed or adaptive mesh refinement.

The field of NR has made significant progress in previous decades, particularly with the development of powerful numerical techniques and advances in high-performance computing. There now exists several catalogs of BBH merger simulations [68, 69] which are frequently used (often in conjunction with other waveform modelling techniques) for GW data analysis and full IMR waveform modelling. However, despite its successes, the main production codes have not changed their numerical methods in the last 10-15 years, and NR still faces several challenges. One of the main challenges is the computational cost of simulations, which can be prohibitively expensive for lower mass binaries with longer waveforms in the frequency band observed by LIGO-Virgo-KAGRA. Therefore to produce longer waveforms which extend far into the inspiral for BBH signals, and cover a significant portion of the evolution of binaries containing neutron stars, NR waveforms often require hybridisation with PN or other semianalytic approximations to extend their length. We will discuss approaches to combine PN and NR information to produce full IMR waveforms in Chapter 3. Additionally, practical difficulties in formulation and convergence of simulations for more complex binaries, such as moderate-high mass ratio ($q \gtrsim 18$) and highly spinning and precessing binaries, means that parameter space coverage of NR simulations in these areas of parameter space remains sparse. Therefore semianalytic approaches to waveform modelling are still a crucial tool for GW data analysis to bridge these gaps. In the following section, we discuss BH spin and precession further, highlighting their effect upon a gravitational waveform and the resulting added complexity of a waveform from a binary which includes precession.

2.4 Effects of Black Hole Spin and Precession

In this section, we discuss in more detail the effect of BH spin upon the gravitational waveform. We separate these effects into two distinct scenarios, firstly when the BH spins

are aligned with the orbital angular momentum \vec{L} , and secondly when they are misaligned and display the spin precession. The spin of a BBHs is an information tracer of the binary's formation history, and accurate measurements of BH spins will help constrain BBH formation channels [21–29, 31–36]. For a more thorough review of BBH formation channels, we refer the reader to, eg. Refs. [21, 70]. In this section, we describe the behaviour of both spin-aligned and precessing binaries, and the resulting waveform phenomenology.

2.4.1 Aligned Spin Binaries

If the binary is non-spinning, $\vec{\chi}_1 = \vec{\chi}_2 = 0$, or the spins are aligned with the orbital angular momentum, $\hat{L} \cdot \vec{\chi}_1 = \hat{L} \cdot \vec{\chi}_2 = 0^1$, then the orbital plane remains fixed in time and the trajectory of the BHs will be confined to this fixed plane through inspiral to merger. Therefore, if the BH vectors χ_1, χ_2 are (anti-)aligned with \hat{L} at some initial reference time t_0 , ie. if we have

$$\chi_{1x}(t_0) = \chi_{1y}(t_0) = \chi_{2x}(t_0) = \chi_{2y}(t_0) = 0, \quad (2.92)$$

then they will remain (anti-)aligned. The predominant effect of the aligned spins is to affect the inspiral rate, and this influence can be parameterised by a weighted sum of the two spins χ_{1z}, χ_{2z} , and so efficient waveform models for aligned-spin binaries can be produced with only one spin parameter rather than two [71–75]. This parameter χ_{eff} is defined as

$$\chi_{\text{eff}} = \frac{m_1 \chi_{1z} + m_2 \chi_{2z}}{M} \quad (2.93)$$

and is a constant of the 2PN equations of motion in the absence of radiation reaction [76], and approximately constant when radiation reaction is included. For example, the PN expression for the binary phase up to 1.5PN order is given by,

$$\Phi(v) = \Phi_0 - \frac{1}{32v^5\eta} \left(1 + v^2 \left[\frac{55\eta}{12} + \frac{3715}{1008} \right] + v^3 \left[\frac{565}{24} \chi_{\text{PN}} - 10\pi \right] \right), \quad (2.94)$$

¹We note the difference between the orbital angular momentum \hat{L} and its Newtonian approximation \hat{L}_N ; however, in this thesis we will not distinguish between them.

where

$$\chi_{\text{PN}} = \chi_{\text{eff}} - \frac{38\eta}{113}(\chi_{1z} + \chi_{2z}). \quad (2.95)$$

Therefore if χ_{eff} is positive, the phase evolution will be slower, the inspiral rate will decrease, and the inspiral will be longer. Similarly if χ_{eff} is negative, the inspiral will be shorter. This is known as the orbital hangup effect [263]. Additionally, the magnitudes of the spins will remain approximately constant throughout the inspiral, therefore the intrinsic parameter space for aligned spin binaries (scaling out the total mass M) is 3-dimensional in mass ratio q , χ_{1z} and χ_{2z} . In terms of waveform phenomenology, the GW signal from a spin-aligned binary will be qualitatively very similar to that of a non-spinning binary (for example as shown in Fig. 2.2), with the amplitude and phase increasing monotonically in the inspiral up to merger, but a modified rate of phasing due to the spins.

We briefly note that as well as the degeneracy between aligned binary spins which map to the same χ_{eff} , there also exists for aligned-spin binaries a parameter degeneracy between mass ratio q and χ_{eff} [264]. The both of these parameters have an effect of modifying the inspiral rate.

2.4.2 Spin Precession

If the black holes' spins are misaligned with the direction of the orbital angular momentum \hat{L} of the binary motion, spin-induced precession of the orbital plane occurs [14, 15, 77, 78], breaking the equatorial symmetry of the system. The precession of the orbital plane is driven by the spin components $\vec{\chi}_{1\perp}$ and $\vec{\chi}_{2\perp}$ instantaneously perpendicular to \hat{L} , defined as $\vec{\chi}_{i\perp} = \vec{\chi}_i \times \hat{L}$. This can be seen from the equations of motion for a precessing binary,

which at 2PN order, and averaged over a single circular orbit, are given by [14]:

$$\begin{aligned} \dot{\hat{L}} = & \frac{L}{R^3} \left[\left(2 + \frac{3m_2}{2m_1} \right) \vec{S}_1 + \left(2 + \frac{3m_1}{2m_2} \right) \vec{S}_2 \right] \times \hat{L} - \frac{3}{2} \frac{1}{R^3} \left[(\vec{S}_2 \cdot \hat{L}) \vec{S}_1 + (\vec{S}_1 \cdot \hat{L}) \vec{S}_2 \right] \times \hat{L} \\ & - \frac{32}{5} \frac{\mu^2}{R} \left(\frac{M}{R} \right)^{5/2} \hat{L}, \end{aligned} \quad (2.96)$$

$$\dot{\vec{S}}_1 = \frac{1}{R^3} \left[\mu M^{1/2} r^{1/2} \left(2 + \frac{3m_2}{2m_1} \right) \hat{L} \right] \times \vec{S}_1 + \frac{1}{R^3} \left[\frac{1}{2} \vec{S}_2 - \frac{3}{2} (\vec{S}_1 \cdot \hat{L}) \hat{L} \right] \times \vec{S}_1, \quad (2.97)$$

$$\dot{\vec{S}}_2 = \frac{1}{R^3} \left[\mu M^{1/2} r^{1/2} \left(2 + \frac{3m_1}{2m_2} \right) \hat{L} \right] \times \vec{S}_2 + \frac{1}{R^3} \left[\frac{1}{2} \vec{S}_1 - \frac{3}{2} (\vec{S}_2 \cdot \hat{L}) \hat{L} \right] \times \vec{S}_2. \quad (2.98)$$

Therefore in a precessing binary, the direction of \hat{L} and the orientation of the two spins \hat{S}_i become time dependent. The precession equations above include the leading order spin-orbit couplings (terms proportional to $\hat{L} \times \vec{S}_i$) and spin-spin couplings (terms proportional to $\vec{S}_1 \times \vec{S}_2$). The spin-orbit couplings are the most dominant effect of precession, occurring at 1.5PN order, and induce the precessional motion of the orbital plane. This orbital plane precession is known as Lense-Thirring precession, and is a purely relativistic effect not present in Newtonian gravity, as are the higher PN order spin coupling effects. The spin-spin couplings start at 2PN order, and introduce nutation of the spins and orbital angular momentum. We note, however, that spin nutation has also been shown to occur in the case of single spin precessing binaries.

In the case of simple precession (which covers all except a few specific spins, and upon which we will focus exclusively hereafter), throughout the inspiral $\hat{L}(t)$ traces a cone centered around the direction of the total angular momentum \hat{J} , which remains approximately spatially fixed [14, 15], i.e., $\hat{J}(t) \simeq \hat{J}_{t \rightarrow -\infty}$, where $\vec{J} = \vec{L} + \vec{S}_1 + \vec{S}_2$. We note that this fixed direction can be seen as a direct consequence of the precession equations at 2PN order (Eqs. (2.96)-(2.98)), and so the direction of \hat{J} is fixed at this PN order. The opening angle of this precession cone $\lambda_L(t)$ is defined as [14]

$$\cos(\lambda_L(t)) \equiv \hat{L}(t) \cdot \hat{J}(t) = \frac{L(t) + S_{\parallel}(t)}{\sqrt{(L(t) + S_{\parallel}(t))^2 + S_{\perp}^2(t)}}, \quad (2.99)$$

where $\vec{S}(t) = \vec{S}_1(t) + \vec{S}_2(t)$ is the total spin of the binary with $S_{\perp}(t) = |\vec{S}(t) \times \hat{L}(t)|$ and $S_{\parallel}(t) = \vec{S}(t) \cdot \hat{L}(t)$. The precession cone opening angle grows on the precession timescale, which lies between the shorter orbital timescale and the longer radiation reaction timescale, ie. the time it takes for the binary to merge.

Precession introduces characteristic amplitude and phase modulations in the GW signal, as can be seen in Fig. 2.4. In this Figure, we see the polarisations $h_+(t)$ for a precessing binary with parameters $\{M, q, \chi_1, \chi_2, d_L, \iota, \psi, \Phi_0\} = \{100M_{\odot}, 3, \{0.6, 0.3, 0\}, \{0.1, -0.9, 0\}, 100\text{Mpc}, 0, 0, 0\}$, where Φ_0 is quoted at a reference frequency of 4Hz. The modulations in the amplitude due to spin precession can be seen clearly, and the amplitude no longer increases monotonically with time as in the case of a spin-aligned or non-spinning binary. Precession also induces phase modulations into the GW signal, though these are not visible in the Figure presented here. We can also clearly see the separation of timescales, as the time taken for the binary to complete a single orbital (double the timescale of one complete GW oscillation), is much shorter than the precessional timescale over which the amplitude modulations vary. This in turn is much shorter than the overall timescale for the binary to inspiral, the radiation-reaction timescale.

Precession also introduces a richer mode structure into the GW signal as will be discussed in more detail in Sec. 3.2, and modifies the properties of the merger remnant object from those of a spin-aligned binary, such as the final spin and recoil velocity of the remnant BH.

Due to their complexity, precessing waveforms encode vast amounts of information which can be used to break parameter degeneracies [16–19], such as the aforementioned mass-spin degeneracy between the mass ratio and χ_{eff} [79–81] as one example.




Figure 2.4: Polarizations h_+ , h_x of a precessing BBH coalescence with parameters $\{M, q, \vec{\chi}_1, \vec{\chi}_2, d_L, \iota, \psi, \Phi_0\} = \{100M_\odot, 3, [0.6, 0.3, 0], [0.1, -0.9, 0], 100 \text{ Mpc}, 0, 0, 0\}$ (where Φ_0 , $\vec{\chi}_1$ and $\vec{\chi}_2$ are quoted at a reference frequency of 4Hz), showing the induced amplitude modulations. The waveforms were produced with the waveform model SEOBNRv4PHM [48] within LALSimulation [49].

2.5 Astrophysical Inference from Gravitational Wave Detections

In Sec. 2.2.2 we described how the GW signal from a merging binary depends on its parameters, and orientation relative to a GW detector. However, we did not consider how, given data from a detected GW signal, we may actually go about determining those binary parameters. In this section, we outline the main concepts behind parameter estimation for BBH signals. The fundamental ideas presented here will be used as part of our analysis in later chapters, notably as part of the inference performed in Chapter 6. For a more thorough overview of Bayesian inference in GW astronomy we refer the reader to Ref. [82].

The aim of BBH parameter estimation is, given detector strain data d (see Eq. (2.64)), to construct a posterior probability density function for the BBH parameters $\vec{\theta}$ which

produced that strain data,

$$p(\vec{\theta}|d), \quad (2.100)$$

which is a normalised such that

$$\int d\vec{\theta} p(\vec{\theta}|d) = 1. \quad (2.101)$$

We begin with Bayes' theorem,

$$p(\vec{\theta}|d) = \frac{\mathcal{L}(d|\vec{\theta}) \pi(\vec{\theta})}{\int d\vec{\theta} \mathcal{L}(d|\vec{\theta}) \pi(\vec{\theta})}, \quad (2.102)$$

where \mathcal{L} is the likelihood function of the data d given the BBH parameters $\vec{\theta}$, $\pi(\vec{\theta})$ is our prior distribution on $\vec{\theta}$, and the denominator is the evidence, which here acts as a normalisation factor. The likelihood describes the probability of a particular detector output, given the BBH parameters $\vec{\theta}$ which describe the astrophysical signal. We may choose the functional form of the likelihood function, and by doing so we implicitly introduce a model for the noise. For gravitational wave detectors, we usually choose the following form:

$$\mathcal{L}(d|\vec{\theta}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2} \frac{|d - \mu(\vec{\theta})|^2}{\sigma^2}\right), \quad (2.103)$$

and assume that the detector noise σ is stationary over the timescale of a CBC signal, and Gaussian with zero mean. Here π refers to the mathematical constant, and $\mu(\vec{\theta})$ is a representation of the gravitational wave strain given $\vec{\theta}$.

Similarly to the likelihood function, the prior $\pi(\vec{\theta})$ is also something we may choose, and represents our prior belief about the parameters $\vec{\theta}$. If we are ignorant about $\vec{\theta}$, we may choose priors such as a uniform distribution which has support across a broad range of parameter space, for example. We note that since a quasicircular BBH signal is parameterised by 15 parameters, the prior $\pi(\vec{\theta})$ and posterior $p(\vec{\theta}|d)$ will be 15-dimensional functions. If we wish to concentrate on the posterior distribution of a single parameter

$\theta_i \in \vec{\theta}$, we may marginalise over other parameters to obtain a marginalised posterior

$$p(\theta_i|d) = \int \left(\prod_{k \neq i} d\theta_k \right) p(\vec{\theta}|d) = \frac{\mathcal{L}(d|\theta_i)\pi(\theta_i)}{\int d\vec{\theta} \mathcal{L}(d|\vec{\theta}) \pi(\vec{\theta})}, \quad (2.104)$$

where $\mathcal{L}(d|\theta_i)$ is the marginalised likelihood which we can express as

$$\mathcal{L}(d|\theta_i) = \int \left(\prod_{k \neq i} d\theta_k \right) \pi(\theta_k) \mathcal{L}(d|\vec{\theta}), \quad (2.105)$$

where θ_k are the 14 parameters in $\vec{\theta}$ which are not θ_i .

Waveform models for BBH coalescences, which are a main focus of this thesis, efficiently give information about what the data d should look like given $\vec{\theta}$ (they are used in construction of $\mu(\vec{\theta})$ in Eq. (2.103)). However, the inverse problem in Eq. (2.102) of constructing the posterior distribution is more computationally challenging due to the high dimensionality of 15 BBH parameters to constrain. To tackle this problem, typically in GW data analysis we make use of stochastic samplers, which generate posterior samples $\{\vec{\theta}\}$ drawn from the posterior distribution $p(\vec{\theta}|d)$ such that the number of samples in some small interval $(\vec{\theta}, \vec{\theta} + \Delta\vec{\theta})$ is proportional to the posterior probability $p(\vec{\theta}|d)$. Two of the most commonly used algorithmic approaches to sampling for GW inference are nested sampling [83] and Markov-chain Monte Carlo (MCMC) sampling [84, 85].

MCMC sampling can be visualised as a set of walkers taking a random walk through the posterior distribution, where the probability of moving to a given point is the transition probability of the Markov chain. Nested sampling is designed to compute the evidence rather than directly draw samples from the posterior distribution, but produces posterior probabilities as a byproduct. For our Bayesian inference in Chap. 6 we make use of the nested sampler *Dynesty* [86] implemented within the *Bilby* [87] Bayesian inference library. We also use the standard priors for gravitational wave data analysis, and refer the reader to [87] for more details.

It is often useful to construct credible intervals from posterior samples as a measure of the uncertainty, and we will make use of the symmetric 90% credible interval. The

90% credible interval on a parameter θ_i is between $\theta_{5\%}$ and $\theta_{95\%}$, where the probability that θ_i lies between $\theta_{5\%}$ and $\theta_{95\%}$ is 0.9. Additionally, the probability that θ_i is less than $\theta_{5\%}$ is 0.05, and the probability θ_i is less than $\theta_{95\%}$ is 0.95. We also note here that the maximum likelihood point (MaxL) or maximum posterior probability point (MaP) are often used as the 'best-guess' recovered parameter values. As their names suggest, the MaxL point corresponds to the posterior sample values of $\vec{\theta}$ for which the likelihood $\mathcal{L}(d|\vec{\theta})$ is maximised, and the MaP is the posterior sample values of $\vec{\theta}$ which have the maximum posterior probability $p(\vec{\theta}|d)$.

As larger number of GW events from CBC mergers are detected, we may begin to think of these as not only individual event detections, but in terms of a population of black holes and neutron stars. Therefore we may use information from individual events, and the posteriors on the parameters of the individual binary masses, spins and other parameters, to constrain the distribution of these parameters across the population. This concept of population inference is discussed in more detail in Chap. 6, and the mathematics of hierarchical Bayesian inference, which is the tool we use to constrain the population distribution using individual events, is introduced in Sec. 6.1

2.6 Summary

In this chapter, we have introduced Einstein's field equations as our best description of gravity in Sec. 2.1, deriving the expressions for GWs first in linearised gravity first in the vacuum in Subsec. 2.1.1, and then including source information in Subsec. 2.1.2. We then discussed astrophysical sources of GWs, introducing CBCs in Subsec. 2.2.1 and then describing how the parameters of the binary and their orientation relative to a ground-based detector affect their phenomenology in Subsec. 2.2.2. We then briefly introduced some other astrophysical sources of GWs in Subsec. 2.2.3. We then discussed the approaches to modelling GW signals from BBH mergers, with the PN formalism for the weak field in Subsec. 2.3.1 and the strong field and NR in Subsec. 2.3.2. We explicitly discussed

the effect of BH spin, which forms a main focus of this thesis, both aligned spins in Subsec. 2.4.1 and precession in Subsec. 2.4.2. Finally, we introduced the main concepts behind Bayesian inference for parameter estimation of BBH sources in Sec. 2.5.

It is crucial to note that each likelihood evaluation in the sampling process requires the evaluation of a gravitational waveform model, which takes as inputs the binary parameters $\vec{\theta}$ and outputs a corresponding GW signal. The inference of the source parameters $\vec{\theta}$ in an unbiased way is therefore predicated on the availability of accurate theoretical models of the emitted GW signal through inspiral, merger and ringdown (IMR). Additionally, on average $10^6 - 10^8$ such model evaluations are needed for parameter estimation of a single GW event to obtain well-sampled posterior probability distributions, and so fast model evaluation speeds are a necessary requirement.

Therefore, in Chapter 3 we will introduce waveform modelling and current approaches in more detail, with a specific focus on waveform models which include precession.

3. Waveform Modelling for Scientific Interpretation of Gravitational Waves

This chapter includes portions of reformatted text from Ref. [40], "Accelerating Multimodal Gravitational Waveforms from Precessing Compact Binaries with Artificial Neural Networks", by Thomas et al. Notably, the introduction to surrogate modelling in Sec. 3.5, as well as small portions of text in Secs. 3.1 and 3.6.

As described in Sec. 2.5, the extraction of the binary parameters from CBC GW signals, and therefore the potential to infer astrophysical information from these signals, is heavily dependent upon theoretical and computational models of gravitational waveforms. In Sec. 2.3.1 we saw how in the weak field, a PN approximation can be used to give the equations for the motion of the binary, and therefore the gravitational wave signal, as an expansion in powers of the characteristic binary velocity when it is sufficiently slow-moving $v/c \ll 1$. Then in Sec. 2.3.2, we noted that in the strong field regime, the Einstein equations must be solved numerically for the binary motion and GW emission. Therefore in order to make efficient use of GW data from CBCs, current waveform models combine information from both NR and analytical approximations to model the full inspiral, merger and ringdown.

The current generation of IMR waveform models can be broadly divided into three families: phenomenological (phenom) models [71, 72]; Effective-One-Body (EOB) models [88]; and surrogate models [89], including NR surrogate models [90, 91]. In this chapter, we give an overview of the main approach underlying each of these waveform families, with particular attention paid to how the effects of spin precession are modelled. In Sec. 3.1, we first describe how the full gravitational wave strain is generally decomposed into the more easily-modelled components of waveform mode amplitudes and phases. Then in Sec. 3.2, we describe in more detail how the effects of precession are imprinted on the GW signals, and the approach used by current waveform models of decomposing

the fully precessing signal into a simpler waveform in a time-dependent reference frame. Finally, in Sec.s 3.4, 3.3, and 3.5, we briefly introduce the underlying frameworks behind each of the three main waveform model families (phenom, EOB and surrogate models), specifically mentioning the main features of these models in terms of speed, accuracy, and their treatment of precession.

3.1 Multipole Waveform Decomposition

We define the complex strain $h(t)$ as the combination of the gravitational wave polarisations:

$$h(t) = h_+(t) - ih_\times(t), \quad (3.1)$$

where the polarisations h_+ , h_\times are defined in Eq. 2.64. We note that this strain is distinct from the detector strain of Eq. 2.64 as it does not include the detector beam pattern coefficients. For waveform modelling purposes these coefficients are unimportant as they encode only the relative orientation of the detector with respect to the source, rather than the intrinsic physics of the binary. Similarly, it is convenient for modelling purposes to decompose $h(t)$ into GW modes in terms of spin-weighted spherical harmonics with spin weight $s = -2$, which encode the directional dependence of the gravitational radiation field. We introduce basis functions on the unit sphere:

$$Y_{\ell m}^{-s}(\theta, \phi) = (-1)^s \sqrt{\frac{2\ell + 1}{4\pi}} d_{ms}^\ell(\theta) e^{im\phi}, \quad (3.2)$$

where s is the spin weight, d_{ms}^ℓ denotes the small-d Wigner matrices [92] and (θ, ϕ) are the polar and azimuthal angles on the sphere. The explicit values for $Y_{\ell m}^{-2}$ can be found in App. D of Ref. [93], or equivalently App. A of Ref. [94]. With this basis decomposition, our strain is now

$$h(t, \theta, \phi; \vec{\theta}) = \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell m}(t; \vec{\theta}) Y_{\ell m}^{-2}(\theta, \phi), \quad (3.3)$$

where we have explicitly noted the dependence of the waveform modes $h_{\ell m}$ on the intrinsic source binary parameters $\vec{\theta}$ (see Eq. 2.69 for details of these intrinsic parameters), which we note are distinct from the polar angle θ . The individual GW modes $h_{\ell m}$ may be extracted by the surface integral

$$h_{\ell m}(t; \vec{\theta}) = \int d\Omega h(t, \theta, \phi; \vec{\theta}) Y_{\ell m}^{-s*}(\theta, \phi), \quad (3.4)$$

where $d\Omega = \sin(\theta)d\theta d\phi$ is the surface element on the unit sphere and $*$ denotes the complex conjugate. As mentioned in Sec. 2.2.2, it is often useful to work in the Fourier domain for GW data analysis, and so we define the Fourier transform of $h_{\ell m}$ in accordance with Eq. 2.73 as

$$\tilde{h}_{\ell m}(f; \vec{\theta}) = \int_{-\infty}^{\infty} h_{\ell m}(t; \vec{\theta}) e^{2\pi i f t} dt. \quad (3.5)$$

Each of these modes may then be further decomposed into a mode amplitude $A_{\ell m}$ and phase $\phi_{\ell m}$, and so the GW modes given by

$$h_{\ell m}(t; \vec{\theta}) = A_{\ell m}(t; \vec{\theta}) e^{-i\phi_{\ell m}(t; \vec{\theta})}, \quad (3.6)$$

or equivalently in the frequency domain,

$$\tilde{h}_{\ell m}(f; \vec{\theta}) = \tilde{A}_{\ell m}(f; \vec{\theta}) e^{-i\tilde{\phi}_{\ell m}(f; \vec{\theta})}, \quad (3.7)$$

where $h_{\ell m}(t; \vec{\theta})$ and $\tilde{h}_{\ell m}(f; \vec{\theta})$ are related to each other through Eq. 3.5. Waveform models generally split the full signal into modes which are then modelled separately, and a common approach for some models (notably Phenom and surrogate models) is to further break these modes into amplitudes and phases. The modes where $\ell = 2$ and $|m| = \pm 2$ are the quadrupole modes, and are usually the dominant harmonics (ie. the largest amplitude and so largest SNR). Additionally, for non-precessing systems the equatorial symmetry

about the orbital plane of the binary implies that

$$h_{\ell m}(t) = h_{\ell -m}^*(t). \quad (3.8)$$

Therefore historically, waveform models tended to include only the dominant (2, 2)-mode, and neglect the subdominant effect of higher (ℓ, m) -modes [265–268]. Fig. 3.1 shows the amplitudes of the Fourier domain modes $\tilde{h}_{\ell m}$ for a non-spinning, equal-mass binary of total mass $50M_{\odot}$ as solid lines, with an inclination of $\iota = 0$. The waveform modes were created using the IMRPhenomXPHM waveform model [168] as implemented in LALSimulation. It shows the hierarchy of mode power, with the (2, 2)-mode dominant in the inspiral, and because there is no mass or spin asymmetry there is little power in the odd- m modes, the (2, 1) and (3, 3)-modes. ¹ In dashed lines we see the same binary with a mass ratio of $q = 2$, and while the amplitudes of the even- m modes ((2, 2), (3, 2) and (4, 4)) remain almost unchanged, the (2, 1) and (3, 3) have much more power, showing that the mass asymmetry has excited these odd- m modes. We note that spin asymmetries will have a similar effect.

Neglecting higher modes in waveform models can lead to significant loss of detection rate [95–98] for high mass or high mass-ratio binaries, and systematic bias in the inferred source parameters [95, 99–102]. One reason for this is that binaries with unequal masses can have significant power in their higher modes, as seen in Fig. 3.1, but a further explanation is as follows. We consider a scenario in which the line of sight of an observer to a GW binary source lies along the orbital angular momentum \hat{L} , so $\theta = 0$ in Eq. 3.3. In this case the strain observed in the interferometer will be predominantly comprised of energy from the (2, 2)-mode, ie. Eq. 3.3 becomes

$$h(t, 0, \phi; \vec{\theta}) \approx h_{22}(t, \vec{\theta}) Y_{22}^{-2}(\theta, \phi) = \frac{1}{8} \sqrt{\frac{5}{\pi}} (1 - \cos \theta)^2 e^{-2i\phi} h_{22}(t, \vec{\theta}), \quad (3.9)$$

¹Strictly mathematically there should be zero power in odd- m modes with no mass or spin asymmetries, the small amount of power shown in these modes in Fig. 3.1 is as a result of waveform model inaccuracies.




Figure 3.1: Amplitudes in the frequency domain for the spherical harmonic modes $\tilde{h}_{\ell,m}(f)$ for a non-spinning BBH with total mass $50M_{\odot}$, inclination $\iota = 0$, with a mass ratio of $q = 1$ (solid lines), and $q = 2$ (dashed lines). The waveform modes were created using the `IMRPhenomXPHM` waveform model [168] as implemented in `LALSimulation`. In the inspiral, assuming a face-on (or face-off) orientation, the quadrupolar $(2,2)$ -mode contains the most power and is dominant for both binaries, but the mass asymmetry excites the odd- m modes, which have more power for the $q = 2$ binary.

and so the system is well approximated by only the quadrupolar mode. Because the $(2,2)$ -mode is dominant for non-precessing binaries, it contains the most power and so binaries in this configuration can be observed out to further distances [103, 104]: they are said to be optimally-oriented. If instead the binary is orientated with some inclination ι with respect to the line of sight of the observer, then the interferometer will not observe the full power contained within the quadrupolar mode, and will record a signal where a larger proportion of the power is contained within higher modes. In this case, the accuracy of the approximation to neglect higher modes breaks down. Additionally, Ref. [105] demonstrates that even for face-on binaries, where the signal contains minimal HM content, including HMs in the model improves the inclination measurement, as face-off inclinations can be excluded. Therefore, by including higher modes within waveform models and measuring their relative power, this allows for breaking of the degeneracy between distance and inclination [44, 106–108]. An example of this is shown in Fig. 3.2, where the




Figure 3.2: The posterior distribution for the luminosity distance d_L and inclination θ_{JN} for the gravitational wave event GW190412. The 90% credible regions are shown, in blue for a waveform model that does not include either higher modes or precession, SEOBv4_ROM [110]. The degeneracy can clearly be seen, and the posterior support for θ_{JN} spans the whole range from 0 to π . This degeneracy is broken when higher modes are included using the model SEOBv4HM_ROM [111, 112] in the dashed red contour, and is improved further when precession is included as well with SEOBv4PHM [48] in yellow. Figure taken from Ref. [109].

posteriors for luminosity distance d_L and inclination θ_{JN} are shown for a detected GW signal GW190412 [109], which had a mass ratio of around $q \sim 3.8$. We note that the degeneracy between the inferred distance and inclination in the blue contour, when using a model (SEOBv4_ROM [110]) which does not include either precession or the effect of higher modes, is broken when higher modes are included in the dashed red contour (the model used here is SEOBv4HM_ROM [111, 112]). This measurement is further improved with the inclusion of precession using the model SEOBv4PHM [48], and the yellow contour shrinks even further.

If the binary is precessing, as described in Sec. 2.4.2 this leads to tilting of the orbital plane, which we can think of as a time-dependent binary orientation with respect to

the observer’s line of sight. Similarly to a non-precessing binary which is non-optimally oriented, there will be significant power in the higher modes observed in the GW signal. Additionally, spin asymmetries can be very pronounced for precessing binaries, and so odd- m modes are likely to contain significant power. This complexity of precessing waveforms means they are able to break parameter degeneracies [16–19, 79–81, 113–117], allowing us to more effectively extract astrophysical information, for which accurate modelling of the effects of both precession and higher modes is critical. However, we note that this added complexity will necessarily lead to more complex and therefore slower waveform models, and so these competing effects must be balanced to achieve both complete and timely inference. We discuss in more detail strategies to speed up precessing waveform models, namely using surrogate models and neural networks, in Chapter 5.

3.2 Modelling Precession

In this section we summarise the approach and assumptions used to model precessional effects in full IMR waveform models. For a more detailed discussion of precessing phenomenology we refer the reader to Refs. [118, 119], Refs. [14, 15] for a comprehensive treatment of precessing binaries within the PN framework, and to Ref. [120] for a recent analysis of precessing modelling assumptions and their accuracy impact.

In Sec. 2.4.2 we describe how the effects of precession create richer morphology and structure of gravitational waveforms, complicating efforts to model such signals. We also describe the separation of the different orbital and precessional timescales in the inspiral, which allows us to define approximate closed-form solutions of the post-Newtonian (PN) precession equations in the inspiral [121, 122]. This separation of timescales motivates that the dynamics of the binary orbit and the precession dynamics can also be approximately separated in the inspiral. It was first suggested in Ref. [14] that the total phase of a precessing signal could be described by modulation of an underlying carrier phase. In a non-precessing binary the time derivative of the orbital phase $\dot{\phi}$ is equal to the angular

3. WAVEFORM MODELLING

velocity of the two bodies in the orbital plane Ω . However, in a precessing binary where the orbital plane becomes time-dependent this no longer holds, and so we can define a carrier phase given by

$$\phi_C(t) = \int_{t_0}^t \Omega(t') dt', \quad (3.10)$$

where t_0 is some (arbitrary) initial time. In this construction, the precessing GW strain $h(t)$ can then be modelled as

$$h(t) = \Lambda(t) h_C(t), \quad (3.11)$$

where $h_C(t)$ is a "carrier signal" of a non-spinning binary with phase specified by Eq. 3.10, and $\Lambda(t)$ is a factor which represents the precession-induced modulations of the GW amplitude and phase. To improve agreement between true precessing waveforms and these artificially modulated constructions, Ref. [123] reformulated $\Lambda(t)$ to better capture the precession-induced modulations, though this came at the cost of introducing several non-physical parameters which could be shown to mimic detector noise and decrease search sensitivity [124].

A key breakthrough in modelling precession effects came in Ref. [173] with the realisation that a fully precessing signal may be greatly simplified by a coordinate transform into a time-dependent reference frame which approximately tracks the precession of the orbital plane. In this non-inertial frame, called the coprecessing frame, the mode hierarchy of a non-precessing but aligned-spin binary is approximately restored, which motivates the definition of a mapping between a fully precessing binary characterised by seven intrinsic parameters, and a corresponding appropriate aligned-spin binary which is fully described by three parameters (see Eq. 2.69) [128]. Therefore the complex problem of modelling a fully precessing waveform is broken down into two more manageable pieces: (i) modelling a simplified waveform in a coordinate frame which tracks the precession of the orbital plane; and (ii) modelling the time-dependent rotation that transforms the signal in this

coordinate frame back to an inertial frame. We can express this quantitatively as

$$h_{\ell m}^P(t) \simeq \sum_{m'=-\ell}^{\ell} \mathbf{R}_{\ell m m'}(t) h_{\ell m'}^{\text{coprec}}(t), \quad (3.12)$$

where h_{ℓ}^P is the fully precessing signal in an inertial frame, $h_{\ell m'}^{\text{coprec}}(t)$ is the simplified waveform, and $\mathbf{R}_{\ell m m'}(t)$ is the time-dependent rotation from the coprecessing frame to the inertial frame.

A common way to represent this rotation in waveform models is to use Euler angles, $\{\alpha(t), \beta(t), \gamma(t)\}$, which represent the time-dependent rotation from a frame in which the z -axis is aligned with the instantaneous orbital angular momentum $\hat{L}(t)$ (henceforth called the L -frame), and an inertial frame (the J -frame) in which the z' -axis is aligned with the total angular momentum \hat{J} evaluated at some reference time, though as described in Sec. 2.4.2 the direction of J remains approximately fixed throughout the binary inspiral for simple precession, which makes up the vast majority of precessional morphologies observable in the current ground-based frequency sensitivity range. The rotation described by these Euler angles can be seen in Fig. 3.3, where the L -frame is defined such that the x and y -vectors lie within the orbital plane, the x -axis points from the smaller black hole m_2 to the larger one m_1 , and y points such that $\{\hat{x}, \hat{y}, \hat{z}\}$ forms a right-handed orthonormal basis. The angle γ is defined by the minimal rotation condition [123, 125] $\dot{\gamma} = -\dot{\alpha} \cos \beta$ for uniqueness. We note that this choice of inertial frame is not unique, for example recent EOB models [48, 126] utilise an inertial frame in which the z' -axis is instead aligned with the Newtonian orbital angular momentum L_N at some reference time, but the basic formalism remains the same. The evolution of these angles is usually computed using semianalytic methods for current precessing waveform models [48, 127, 133, 165, 168, 269–271], which we describe in more detail for each waveform family.

As an alternative to Euler angles, quaternions may also be used to represent the time-dependent rotation. If we consider a rotation of angle θ about a unit axis $[x, y, z]$, then this may be described by the quaternion $\hat{q} = [\cos(\theta/2), x \sin(\theta/2), y \sin(\theta/2), z \sin(\theta/2)]$.




Figure 3.3: Definition of the inertial J -frame and the Euler angles. The three Euler angles α , β and γ define the rotation from the L -frame where the z -component is parallel to the orbital angular momentum of the binary $\hat{z} = \hat{L}(t_0)$, to the J -frame where it is parallel to the total angular momentum $\hat{z}' = \hat{J}(t_0)$ at the start time of the waveform t_0 . The angle γ is defined by the minimal rotation condition, $\dot{\gamma} = -\dot{\alpha} \cos \beta$ [123, 125].

Notably, the precessing surrogate NRSur7dq4 [160] uses quaternions to denote the time-dependent rotation from the coprecessing frame to the inertial frame, as described in more detail in Subsec. 4.2.1.

The waveforms in this coprecessing frame resemble those from spin-aligned binaries [128, 129], and so a common approximation in many of the current generation of waveform models is to set

$$h_{\ell m}^{\text{coprec}}(t; q, \{\chi_{1\parallel}, \vec{\chi}_{1\perp}\}, \{\chi_{2\parallel}, \vec{\chi}_{2\perp}\}) = h_{\ell m}^{\text{AS}}(t; q, \chi_{1\parallel}, \chi_{2\parallel}) \quad (3.13)$$

in the RHS of Eq. 3.12, thus "twisting up" an aligned spin waveform, which depends only on the spin components parallel to the orbital angular momentum², to produce a fully precessing one, which depends on all six spin components. We note, however, that this

²We also note that as the presence of precession will modify the spin of the remnant BH, this waveform approximation will also depend on the final spin $\vec{\chi}_{\text{remnant}}$.

approximation introduces significant systematic errors due to the neglect of spin-induced mode asymmetries in precessing systems [120, 130], because in a precessing binary the equatorial symmetry of the system is broken and so Eq. 3.8 no longer holds [131].

Therefore, with this simplification we are left with the problem of modelling the precession dynamics, which is complicated by the high dimensionality of precessing systems as it depends on all six spin components. Similarly to the parameter χ_{eff} (Eq. 2.93) described in Subsec. 2.4.1 to model the inspiral rate, it is useful to employ strategies of dimensional reduction to model the effects of precession with fewer degrees of freedom. Motivated by PN theory, Ref. [174] introduced an effective precession spin χ_p defined as

$$S_p := \max(A_1 S_{1\perp}, A_2 S_{2\perp}), \quad (3.14)$$

$$\chi_p := \frac{S_p}{A_1 m_1^2}, \quad (3.15)$$

where $A_1 = 2 + 3/2q$, $A_2 = 2 + 3q/2$, and $S_{i\perp} = \|\vec{S}_{i\perp}\|$ such that the Kerr limit $\chi_i \leq 1$ is obeyed. It is motivated by the observation that over the course of the inspiral, the dimensional in-plane spin vectors $\vec{S}_{1\perp}$ and $\vec{S}_{2\perp}$ will rotate at different rates, and so the magnitude of their vector sum will oscillate between the sum and difference of their magnitudes. Therefore χ_p is constructed such that it captures the average amount of precession exhibited by a generically precessing system over many precession cycles defined at some reference time t_{ref} during the inspiral. We note that χ_p will assume a (slightly) different value depending on t_{ref} but this time (frequency) dependence can be mitigated through the inclusion of additional precession-averaged spin effects [132]. An alternative effective parameterisation based on the total spin can be found in [133]. The effective precession spin χ_p is regularly used to make statements about the measurement of precession at a certain reference frequency (time) in GW inference, see e.g. [19, 134, 254], and has been successfully utilised in some Phenom waveform models [127, 269, 272] to model the effects of precession on waveforms, as described in more detail in Subsec. 3.4.4.

In the remainder of this chapter, we briefly outline the modelling framework for each

of the three main IMR waveform model families: EOB in Sec. 3.3; Phenom models in Sec. 3.4; and reduced-order surrogate models in Sec. 3.5. We pay particular attention to how they include the effects of precession and higher multipoles.

3.3 Effective-One-Body Models

A semi-analytic approach to waveform modelling is using the EOB framework [88, 273–276], which combines information from analytical methods including PN theory with NR simulations. In this formalism, the two-body dynamics of a compact binary is mapped onto the dynamics of a test particle moving in an effective metric of deformed Kerr spacetime. EOB waveform models are made up of three key ingredients: (i) the EOB Hamiltonian, (ii) the GW modes, and (iii) the radiation-reaction (RR) force. We will look at each of these pieces in more detail, before describing how they fit together to produce accurate waveforms for CBCs, including the effects of precession.

3.3.1 Effective-One-Body Hamiltonian

The first key piece of the EOB approach is the EOB Hamiltonian, which is defined by the energy map:

$$H_{\text{EOB}} = M \sqrt{1 + 2\nu \left(\frac{H_{\text{eff}}}{\mu} - 1 \right)}, \quad (3.16)$$

where M is the binary’s total mass as defined in Eq. 2.50, and ν is the symmetric mass ratio of Eq. 2.48. The effective Hamiltonian H_{eff} describes the geodesic motion of a test particle of mass μ , where μ is the reduced mass of the binary as defined in Eq. 2.49, in a deformed Kerr spacetime of mass equal to M . We obtain the form of this deformation, which has deformation parameter ν , by requiring that at each PN order, the PN-expanded EOB Hamiltonian agrees with a PN Hamiltonian for BBHs after a canonical transformation, as detailed in Ref. [88]. In the limit of $\nu \rightarrow 0$, the effective dynamics reduces to that of a test particle on a generic orbit around a Kerr black hole. We note that Eq. 3.16 only describes the conservative dynamics of the system, with no damping due to gravita-

tional radiation, and so in this limit the Kerr deformation will be static and spherically symmetric.

There are two possible strategies to incorporate the effects of black hole spin into this formalism, the first of which is to map the spinning dynamics to a test mass in a deformed Kerr spacetime, then including spin contribution corrections into the effective Hamiltonian. The second approach is to map to a test spin in deformed Kerr spacetime, which is a more generic technique as it includes resummed spin-orbit couplings at all PN orders in the test-body limit by construction. However, this generality makes it more complicated than the test mass approach, and so while we note that previous generations of spinning EOB waveform models [48, 110, 111, 126, 166, 267, 268], have used the test spin Hamiltonian method [135, 136], the current state-of-the-art models use both the test mass approach as in Ref. [277], and the test spin approach as in Ref. [278].

3.3.2 Gravitational Waveform Modes

The second key piece to include is the GW modes, which are built upon a factorised resummation of PN modes in the inspiral [137–139]. This resummation extends the validity of the PN approximation to higher velocities, and allows for the use of the EOB mapping approximation up until plunge [273]. These modes in the inspiral-plunge regime are expressed in the form

$$h_{\ell m}^{\text{insp-plunge}} = h_{\ell m} N_{\ell m}, \quad (3.17)$$

where $N_{\ell m}$ are non-quasi circular (NQC) corrections, which allow for better agreement to NR in the late inspiral-plunge stage of the binary dynamics, where the orbit of the evolution is no longer adiabatic. The resummed, factorised PN expressions for $h_{\ell m}$ are currently only derived for non-precessing binaries, and depend on the projections $\chi_{i\parallel}$ of the spins along L . Therefore for precessing binaries where \hat{L} is time-dependent, the EOB dynamics for \hat{L} are used to compute projections $\chi_{i\parallel}(t)$ which are also evolving. In order to obtain the fully precessing waveform modes $h_{\ell m}^P$, the factorised, resummed non-

precessing modes $h_{\ell m}$ which have been evaluated along the EOB precessing dynamics are then "twisted up" into an observer's inertial frame, using Eq. 3.12, using the precessing EOB dynamics.

3.3.3 Radiation-Reaction Force

The third key component of the EOB formalism is the RR force \mathcal{F} , which describes the energy and angular momentum loss due to the emission of GWs. This can be expressed in terms of the GW energy flux in terms of the GW modes as

$$\mathcal{F} = \frac{\Omega}{16\pi} \frac{\vec{p}}{|\vec{L}|} \sum_{\ell=2}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} m^2 |d_L h_{\ell m}|, \quad (3.18)$$

where $\Omega = |\vec{r} \times \dot{\vec{r}}|/|\vec{r}|^2$ is the angular orbital frequency, \vec{r} is the black hole separation, \vec{p} is the canonically conjugate momentum, and d_L is the luminosity distance to the source. We note that these $h_{\ell m}$ are only the factorised, resummed PN modes and not the full inspiral-plunge modes from Eq. 3.17, as these modes do not include NQC corrections, nor do they include the twisting up into an observer's inertial frame for precessing binaries.

3.3.4 Full Formalism

Putting these three key pieces together, the equations of motion in the inspiral-plunge for a BBH with arbitrarily-orientated spins are [276]

$$\dot{\vec{r}} = \frac{\partial H_{\text{EOB}}}{\partial \vec{p}}, \quad (3.19)$$

$$\dot{\vec{p}} = -\frac{\partial H_{\text{EOB}}}{\partial \vec{r}} + \mathcal{F}, \quad (3.20)$$

$$\dot{\vec{S}}_i = \frac{\partial H_{\text{EOB}}}{\partial \vec{S}_i} \times \vec{S}_i + \dot{\vec{S}}_i^{\text{RR}}, \quad (3.21)$$

where H_{EOB} is the EOB Hamiltonian, and \mathcal{F} is the RR force. We note that the RR contribution to the spin-evolution equations (Eq. 3.21) starts at 5.5PN spin-spin order, and can thus be neglected. The model then solves for the EOB dynamics by evolving

these equations in time, thus extracting the resulting GW modes of Eq. 3.17. EOB models are therefore inherently constructed in the time domain, and are computationally more expensive than Phenom models.

The formalism as described thus far is only valid when the BBH is composed of two individual objects, ie. up to merger. Therefore EOB waveforms in the merger-ringdown portion are constructed from functional fits to NR, in a similar spirit to the Phenom approach described below in Subsecs. 3.4.1 and 3.4.2. We also mention that the Hamiltonian mapping described for the inspiral-plunge in Subsec. 3.3.1 is formulated such that there are free calibration parameters which are fitted to NR waveforms. As an example, in recent EOB models [110], the calibration parameters are used to describe the shape of the EOB radial potential, the time delay between peak orbital frequency and peak radiation, as well as the strength of spin-orbit and spin-spin couplings in the effective Hamiltonian. This NR calibration further increases the accuracy of EOB models especially in the aligned spin sector, though we note that currently no EOB model is calibrated to precessing NR.

There currently exists two distinct EOB waveform model families, namely SEOB and TEOB [279–281], which differ in their derivation of the Kerr deformation in the conservative Hamiltonian, as well as treatment of spin corrections (for a recent comparison we refer the reader to Ref. [140]). The most recent generation of SEOB model, SEOBNRv5, uses all information up to 4PN and partial 5PN results in the non-spinning part of the Hamiltonian, and precessing spin contributions up to 4PN as derived in [141]. Both EOB families naturally incorporate precession and higher modes into their Hamiltonian prescriptions, and both make use of the twisting up procedure for the full inspiral-plunge modes as described in Subsec. 3.3.2.

3.4 Phenomenological Models

Another semi-analytic method to waveform modelling for a full IMR signal is the phenomenological approach [71, 72]. Phenom models [105, 127, 165, 167, 168, 265, 266, 269, 270, 272, 282–285] are among the most widely used of waveform models for GW data analysis, as they are computationally extremely efficient. They have been used as part of all GW detections to date, and have been used for data analysis in gravitational wave catalogs such as Refs. [4, 5, 7, 45, 134, 218, 255–257]. They are usually formulated in the frequency domain, and so do not require Fourier transforms for likelihood evaluations, though inherently time-domain Phenom waveform models are presented in Refs. [270, 284, 285]. Phenom models rely on piece-wise closed-form ansatzes to approximate the form of different frequency regions of the waveform, which are usually defined as the inspiral, ringdown, and an intermediate region encompassing the complex physics of the late inspiral and merger. Therefore, the first step is to define appropriate, physically-motivated expressions in each region, which are then fit to a set of calibration waveforms to produce a set of phenomenological coefficients. The coefficients are physically not very meaningful on their own, so a second step is to then map them across the parameter space of BBHs using fits. Finally, the models for the three regions are pieced together, ensuring smoothness of the waveform, to produce a final model for the full IMR signal. Unlike in the EOB approach, Phenom models which are inherently frequency domain do not solve any equations of motion of the binary,³ which makes them significantly cheaper to evaluate but also means that they do not provide the binary dynamics as part of the solution. In the remainder of this section, we briefly outline in more detail the three steps in the Phenom modelling approach in Subsecs. 3.4.1, 3.4.2, and 3.4.3, before describing how the current generation of Phenom models incorporate the effects of precession and higher modes in Subsec. 3.4.4. For a more detailed discussion of the phenomenology of IMR waveforms and how they are modelled in the Phenom family, we refer the reader to,

³The time domain precessing Phenom models `PhenomTP` and `PhenomTPHM` do involve numerical integration of the orbit averaged spin evolution equations, which makes them slower than other precessing Phenom models which are inherently in the frequency domain, but these models do not use the SPA.

eg. Ref. [265].

3.4.1 Closed-Form Expressions and Phenomenological Fits

As we have already mentioned, Phenom models rely on the assumption that if a non-precessing full IMR waveform is broken down into smaller pieces across the full range of frequencies, then each piece can be approximately modelled based on closed-form analytic expressions. The first of these three regions is the inspiral, in which it is natural to model the signal using the PN framework as an expansion in frequency f , where $M\pi f = (v/c)^3$. The most recent generation of Phenom models, **PhenomX** [168, 282, 283], bases this inspiral ansatz on the **TaylorF2** approximant [58, 142, 143, 275], which provides closed-form expressions for the amplitude and phase of the waveform for quasicircular inspirals in the frequency domain using the SPA as described in Sec. 2.3.1. These closed-form expressions are then augmented with higher order, pseudo-PN terms to model behaviour at higher order where the PN coefficients are as yet unknown. More explicitly, the expressions for the inspiral amplitude and phase for the quadrupolar mode are given as in Ref. [282] by

$$A_{\text{inspiral}} = A_0 \left(\sum_{i=0}^6 A_i (\pi f)^{i/3} + \sum_{i=1}^3 \rho_i (\pi f)^{(6i+1)/3} \right), \quad (3.22)$$

$$\phi_{\text{inspiral}} = \phi_{\text{TaylorF2}} + \frac{1}{\eta} \left(\sigma_0 + \sum_{i=1}^5 \frac{3\sigma_i}{2+i} (Mf)^{(2+i)/3} \right), \quad (3.23)$$

where the first term in Eq. 3.22 and ϕ_{TaylorF2} are the analytic amplitude and phase expressions given by the **TaylorF2** approximant, and the coefficients A_i are known. The quantity A_0 factors out the leading order frequency behaviour of the PN expansion of the amplitude, given by $A_0 = \sqrt{\frac{2\eta}{3\pi^{1/3}}} f^{-7/6}$. Therefore the pseudo-PN coefficients ρ_i, σ_i capture the higher order behaviour, which makes a total 9 phenomenological coefficients for the inspiral region. Analogously, in the ringdown region the ansätze for the amplitude and phase derivative are based on Lorentzian functions, as in this region we expect both the amplitude and derivative of the accumulating phase to decay exponentially in

the time domain and so the Lorentzian gives the Fourier transform of this behaviour. In the intermediate regime, the amplitude is based on an inverse polynomial in powers of frequency, and the phase derivative expression is a mixture of polynomial terms and a Lorentzian function, to smoothly match the phase in both the inspiral and ringdown. Full details of these ansätze are given in Ref. [282].

Having defined the closed-form expressions for each region and their corresponding phenomenological coefficients, these expressions are then fitted to a set of calibration waveforms to obtain values for the coefficients. The **PhenomX** generation of models uses waveforms which are a hybridisation of **SEOBNRv4** in the inspiral, matched to NR waveforms at higher frequencies in the late inspiral, merger and ringdown. The hybridisation is performed over a frequency window to ensure smoothness. This calibration set also includes extreme mass ratio waveforms where $q \in \{200, 1000\}$, where the late inspiral, merger and ringdown uses numerical solutions of the perturbative Teukolsky equation [144], and the inspiral uses **SEOBNRv4** as before.

Fitting the phenomenological coefficients to the hybrid waveforms is performed by constructing fits for the values of the amplitudes and phases or at specific frequency nodes, known as collocation points. This provides a system of linear equations for the coefficients which can then be solved to obtain their values.

In previous iterations of Phenom waveform models [105], non-precessing higher modes were modelled by an approximate frequency mapping between the quadrupolar mode and higher modes. The result of this is that the individual modes do not need to be modelled separately, rather the quadrupolar mode can be modelled and then transformed appropriately to represent the higher modes. To increase accuracy, the most recent model, **PhenomXHM**, models each of the modes separately, and so the method above of defining an ansatz and fitting the phenomenological coefficients is performed separately for each mode, amplitude and phase, in each of the three regions.

3.4.2 Modelling Across Parameter Space

Having performed the previous step of fitting for the phenomenological coefficients, these now need to be mapped to the more physically relevant parameters which describe binary black hole mergers. This step is required because any waveform model needs to be able to take in BBH parameters as input, and output a corresponding waveform in order to be useful for inference. This is done in `PhenomX` using a hierarchical fitting procedure as described in Ref. [145], where first the phenomenological coefficients are fitted in η across the one-dimensional subspace of non-spinning binaries. They are then also fitted across the subspace of equal mass, equal spin binaries, and these fits are then combined to produce a 2D ansatz in symmetric mass ratio and spin. Finally, the fit is augmented with information from the spin asymmetries of the binary, to produce a full 3D fit across the parameters for a non-precessing binary.

3.4.3 Full Inspiral-Merger-Ringdown Waveforms

The final stage to generate full non-precessing IMR waveforms from a Phenom model is to glue together the piece-wise closed-form solutions from each of the three regions, to produce a single smooth waveform. This is done in `PhenomXAS` within some frequency matching windows, set relative to the minimum energy circular orbit (MECO), ISCO, and ringdown frequencies to ensure a smooth output waveform.

3.4.4 Including Precession in Phenom Waveforms

The prescription for Phenom waveform modelling as described thus far has been solely based on spin-aligned waveforms. In this section, we outline how Phenom models incorporate the effects of precession. Phenom models, like EOB models, use the approximation of Eq. 3.12 that a non-precessing waveform can be twisted up using time-dependent Euler angles to produce a precessing waveform. In Phenom models, the approximation is made that the simplified waveforms in the coprecessing frame are approximately equivalent to

spin-aligned waveforms. Therefore the final step to producing fully precessing Phenom waveforms is to have a model for the Euler angles, because unlike EOB models, Phenom models do not solve for the binary dynamics and so these angles need to be modelled separately. In the most recent generation of precessing Phenom model, **PhenomXPHM**, there are two possible methods used to model the Euler angles. The first approach uses the next-to-next-to-leading order (NNLO) PN method, as also used in a previous generation of precessing Phenom model **PhenomPv2** [127, 165]. In this approach, the secondary spin is set to zero $\vec{S}_2 = 0$ and thus spin couplings are limited to spin-orbit. The evolution of the Euler angles is expressed in terms of the effective precession spin χ_p [146, 147], which is presumed to be placed upon the primary black hole in this prescription. The second approach, which is the default behaviour for the model, is based upon the application of multiple scale analysis to the PN equations of motion [76, 121, 122, 148–150]. In this framework, the natural separation of timescales between the shorter orbital, longer precessional timescale, and even longer radiation-reaction (or inspiral) timescale, allows for a perturbative expansion approach to solve the precession equations and produce frequency-domain solutions which contain both spin-orbit and spin-spin effects, ie. spin information from both black holes. These Euler angles solutions are given by a PN expansion with an additional MSA correction. It is important to note that by construction, non-spinning and spin-aligned Phenom waveforms are calibrated to numerical relativity, as at their core they are phenomenological fits to NR and NR hybrids. This is no longer true when precession is introduced, as the approximation that precessing waveforms can be built from spin-aligned waveforms twisted up means that precessing NR is not required in their construction, reducing their accuracy in the precessing sector. Until recently, no semi-analytic waveform models (Phenom or EOB) were calibrated to NR in the precessing sector due in large part to the high dimensionality of precessing binaries, which has led to waveforms systematics impeding inference in areas of parameter space where there may be large amounts of precession [38, 218]. The first steps towards precessing NR calibration have been made recently in Ref. [272], in which the aligned-spin Phenom model **PhenomD**

was tuned to the quadrupolar modes of single-spin precessing NR through a simple χ_p parameterisation. However, as previously mentioned in Subsec. 3.2 χ_p alone does not reproduce the higher modes of precessing waveforms, and so a more nuanced strategy may be needed for further development in this area. We discuss this in more detail in Sec. 4.1.

3.5 Surrogate Models

The third and final waveform modelling family we will discuss is that of surrogate waveform models. A surrogate model is a waveform model which is built upon an underlying, slower model, over some predefined parameter space of intrinsic binary parameters. Therefore, their range of validity in parameter space is limited to the region over which they are constructed, the training space, plus an extrapolation region over which the model has been tested and shown to be accurate to within some tolerance. This is in contrast to semianalytic models which have a wide range of parameter space validity are calibrated to NR at high mass ratios and spins. Surrogate models use reduced-order modelling techniques [89, 91, 151, 152] to create a model which is very accurate to the underlying (more expensive) model over that parameter space, but produces waveforms at a fraction of the computational cost. Recent examples of surrogate models for waveforms from coalescing compact binaries include NR [89, 91, 160, 163, 286] and NR-hybrid surrogate models [161], surrogates for the aligned-spin EOB model **SEOBNRv4** [110] using artificial neural networks [153, 154] a machine learning emulation of a different EOB model, **TEOBResumS** [155, 156], as well as other dimensional reduction techniques [157] to greatly reduce the cost of astrophysical inference with these models.

In the following, we will provide a brief outline of the main steps for building a surrogate model. For a more complete explanation we refer the reader to e.g. [152].

The process of building a surrogate model may begin with building a *reduced basis*, which enables us to represent any arbitrary function, e.g. a time-domain waveform $h(t, \vec{\lambda})$ with intrinsic parameters $\vec{\lambda}$, within the discrete training space $\mathcal{T}_M = \{\vec{\lambda}_i\}_{i=1}^M \subseteq \mathcal{T} =$

$\{\vec{\lambda}_i\}_{i=1}^{\infty}$ as a linear combination of an n -dimensional orthonormal basis $\{\hat{e}_i(t)\}_{i=1}^n$ and projection coefficients $\{c_n(\vec{\lambda})\}_{i=1}^n$,

$$h(t, \vec{\lambda}) \approx \sum_{i=1}^n c_i(\vec{\lambda}) \hat{e}_i(t), \quad (3.24)$$

with $n \leq M$. The reduced basis is constructed recursively using a greedy algorithm [151, 152] until all waveforms in the training space \mathcal{T}_M are represented by the basis to within a certain tolerance σ , which is related to the representation error ϵ by

$$\max \epsilon = \max \left\| h(t; \vec{\lambda}) - \sum_{i=1}^n c_i(\vec{\lambda}) \hat{e}_i(t) \right\|^2 \leq \sigma, \quad (3.25)$$

where $\|*\|$ denotes the L^2 -norm, which we compute via the Chebyshev-Gauss quadrature rule. To achieve this, at each step the waveform with the largest representation error using the current basis is chosen, orthogonalised with respect to all current basis elements, and normalised, before being added to the basis as the next basis element. The greedy algorithm stops once Eq (3.25) is fulfilled or if the waveform with the largest representation error is already a basis element. The latter is an indication that the training space \mathcal{T}_M is sampled too coarsely to achieve the desired accuracy σ . If the discrete training space is sampled sufficiently densely, then the reduced basis representation allows us to approximate *any* waveform in the entire parameter space \mathcal{T} .

After the basis has been constructed, we proceed to build an empirical interpolant (EI) using the empirical interpolation method [158, 159], which allows us to reconstruct each waveform $h(t; \vec{\lambda})$ for $\vec{\lambda} \in \mathcal{T}_M$ to within a high accuracy, only using information at certain (sparse) time nodes $\{T_i\}_{i=1}^n$. These carefully selected empirical times or nodes are determined exclusively by the reduced basis waveforms, and the number of time nodes

will be equal to the number of waveforms within the reduced basis:

$$\text{EI}[h](t; \vec{\lambda}) = \sum_{j=1}^n B_j(t) h(T_j; \vec{\lambda}), \quad (3.26)$$

$$= \sum_{i=1}^n \sum_{j=1}^n \hat{e}_i(t) (V^{-1})_{ij} h(T_j; \vec{\lambda}), \quad (3.27)$$

where $(V)_{ij} = (\hat{e}_i(T_j))$ is the interpolation matrix.

The final step for building a surrogate model is to perform a parameter space fit which allows us to predict waveforms at the empirical times $\{T_i\}_{i=1}^n$ for arbitrary parameters $\vec{\lambda} \in \mathcal{T}$ based on the greedy points $\{\vec{\lambda}_i\}_{i=1}^n$ selected to construct the reduced basis. This requires us to fit $h(t; \vec{\lambda})$ across the parameter space at each empirical node such that

$$h(T_i; \vec{\lambda}) \approx A_i(\vec{\lambda}) e^{i\phi_i(\vec{\lambda})}, \quad (3.28)$$

where A_i and ϕ_i are the amplitude and phase at the i -th empirical node. The $2n$ -functions that determine the parameter space fits can be determined by different means, for example via traditional fitting functions such as splines or polynomials [89, 151, 160] or by using machine learning algorithms such as artificial neural networks [40, 153, 154] or Gaussian processes [161–164]. Once the fitting coefficients $A_i(\vec{\lambda})$ and $\phi_i(\vec{\lambda})$ have been determined, the final surrogate model for a waveform $h(t; \vec{\lambda})$ is then given by

$$h^S(t; \vec{\lambda}) \equiv \sum_{i=1}^n \sum_{j=1}^n (V^{-1})_{ji} \hat{e}_j(t) A_i(\vec{\lambda}) e^{-i\phi_i(\vec{\lambda})}. \quad (3.29)$$

We note that this prescription applies to generic functions up to the parameter space fits Eq. (3.28), whose RHS decomposition depends on the function that is being modelled.

3.6 Summary

In this chapter, we have provided a brief introduction to waveform modelling, including the mode decomposition commonly used in waveform models, modelling of precession, and

the main concepts behind each of the three IMR waveform model families. In Sec. 3.1, we described how the waveform strain is decomposed into spherical harmonic modes for modelling purposes, and then in Sec. 3.2 we outlined the key assumptions and approaches for modelling precessional effects in gravitational waveforms. In Secs. 3.3, 3.4, and 3.5 we then discussed the main ingredients of each of the three IMR waveform model families: EOB, Phenom and surrogate models respectively. Recent years have seen significant improvements in the modeling of the complete IMR signal of compact binaries with the inclusion of spin-induced precession effects [165, 166] as well as higher-order harmonics [48, 167, 168]. While the state-of-the-art waveform models are sufficiently accurate for current observations, where the uncertainty in the measurement of the BH properties is dominated by the statistical uncertainty due to detector noise, future upgrades to the current interferometer network [169] and third-generation ground-based detectors such as the Einstein Telescope [219, 224, 287] and Cosmic Explorer [169, 225, 288], will operate at unprecedented sensitivities, shifting focus onto systematic modeling errors as the dominant source of error [170]. The development of evermore accurate models by increasing their physics content is of paramount importance, including the accurate and efficient modelling of effects of precession. Therefore, in the next chapter we discuss calibration of semianalytic waveform models to precessing NR to improve their accuracy, and the need for a dimensional reduction of precessing systems to achieve meaningful calibration of precessing waveforms including higher modes.

4. A New Effective Precession Spin for Modelling Multimodal Gravitational Waveforms in the Strong-Field Regime

This chapter is a reformatted version of work presented in Ref. [39], "New Effective Precession Spin for modelling Multimodal Gravitational Waveforms in the Strong-Field Regime", published in *Phys. Rev. D* 103.83022, for which L. M. Thomas is lead author. The concept for this paper was proposed by P. Schmidt and G. Pratten as a precessing spin dimensional reduction strategy towards calibrating semianalytic precessing waveform models to precessing NR. L. M. Thomas performed the initial investigations for this work to obtain a functional form for χ_{\perp} , with guidance from P. Schmidt and G. Pratten. She also obtained the mode energy thresholds from odd m-modes and sampled the binary parameters for the mismatches, again with input from P. Schmidt and G. Pratten. L. M. Thomas wrote the code for and performed the mode and strain mismatch analyses, and the interpretation of these mismatch results across parameter space was a collaborative effort with L. M. Thomas, P. Schmidt and G. Pratten, including the observation that the improved performance of χ_p is due to better replication of the binary dynamics as opposed to the precession cone opening angle. Therefore, P. Schmidt proposed the idea to perform a mismatch study for the quaternion components and analysis of the merger remnant properties, which were coded up and performed by L. M. Thomas. The final discussion and interpretation of our results was a joint effort between L. M. Thomas, P. Schmidt and G. Pratten. L. M. Thomas produced all the Figs. presented here, and wrote the majority of the text with input from P. Schmidt and G. Pratten.

In Chapter 3, we introduced the main modelling concepts behind each of the three main IMR waveform families currently used for GW data analysis. We outlined how semi-analytic models, both EOB models in Sec. 3.3 and phenom models in Sec. 3.4, are constructed for spin-aligned binaries including calibration to NR in the strong-field to

increase their accuracy. These aligned-spin waveforms are then twisted up using time-dependent rotations to replicate the precessing motion of the orbital plane, and produce a fully precessing waveform. However semi-analytic precessing models do not generally include calibration to precessing NR, in large part due to the high dimensionality of precessing binaries, as described in Sec. 3.2.

In this chapter, we introduce a new effective precession spin $\vec{\chi}_\perp$, which reproduces the quadrupolar and higher mode behaviour of fully precessing binaries with fewer spin components than the full six spin degrees of freedom, effectively reducing the dimensionality of these systems and providing a potential pathway for meaningful calibration to precessing higher mode NR. In Sec. 4.1 we motivate the need for a new effective dimensional reduction of precessing binaries, which includes higher mode behaviour, and then in Sec. 4.1.1 we introduce our new effective precession spin vector $\vec{\chi}_\perp$. We describe the methodology used to assess the efficacy of this dimensional reduction in Sec. 4.2. In Sec. 4.3 we present our results and subsequently discuss the accuracy and caveats of this spin mapping in Sec. 4.4.

4.1 Calibrating Waveform Models to Precessing Numerical Relativity

Semi-analytic models as described in Chapter 3 are constructed to produce gravitational waveforms much faster than by numerically solving the Einstein field equations, using analytic approximations. However, to increase the accuracy of these approximations, these models are tuned to numerical relativity to ensure their agreement across BBH parameter space. For semi-analytic models which do not include the effects of precession, this is now standard practice, with the most recent generation of non-precessing models calibrated to NR across a wide range of mass ratios and aligned spin configurations, see eg. Refs. [277, 283].

Fully tuning aligned-spin models to NR requires calibration in only 3 dimensions:

the mass ratio, and two dimensions of spin, which may correspond to the aligned-spin components χ_{1z} , χ_{2z} or effective combinations thereof. As an example, we consider the recent aligned-spin phenom model **PhenomXHM**, which tunes model parameters to NR across each of the three intrinsic spin-aligned BBH dimensions of mass ratio and spins. These dimensions are explicitly parameterised by the symmetric mass ratio η , a one-dimensional spin parameter \hat{S} whose definition is a function of the effective aligned spin χ_{eff} in the inspiral, and spin asymmetry $\Delta\chi = \chi_1 - \chi_2$. **PhenomXHM** uses a hierarchical fitting technique to calibrate the model parameters to NR across each of these dimensions, for a total of NR 652 simulations. The distribution of these simulations across the parameter space can be seen in Fig. 4.1, though naively if we were to approximate the number of simulations needed for calibration across each dimension to be equal, then the number of NR simulations needed per intrinsic dimension is $\sqrt[3]{652} \sim 10$. We also note that in the example of **PhenomXHM**, due to quality and availability of NR simulations, these calibration points are not distributed evenly across the three dimensions, as shown in Fig. 4.1. Instead, they are mainly clustered around equal mass or non-spinning parts of parameter space, with few points with high mass ratios and high spins.

Current precessing IMR waveform models are built in an approximate way by applying a time-dependent rotation encoding the orbital precession dynamics to waveform modes obtained in a frame that coprecesses with the orbit [128, 173]. This separation of the precessional and orbital dynamics is well-motivated in the early inspiral where the relative timescales are very distinct, but is not expected to be as accurate through merger where the orbital and precessional timescale can become comparable. Indeed, in the currently most widely-used phenom model **PhenomXPHM** the Euler angles are modelled throughout the entire waveform using information from PN theory as described in Sec. 3.4.4, justified by the observation that the PN angles behave smoothly and at high frequencies. While this approximate implementation does give reasonable agreement to NR [165, 167, 168], it is not expected to hold in areas of parameter space with high mass ratios and large in-plane spins. In contrast, recent EOB models [48] make use of an approximation motivated




Figure 4.1: Distribution of the 652 NR simulations used in calibration of PhenomXHM across the three-dimensional intrinsic parameter space for aligned-spin BBHs: mass ratio q ; and aligned spin components χ_1 , χ_2 , which we have referred to as χ_{1z} and χ_{2z} in this thesis. The colours of the dots indicate the code with which the NR simulation was produced, SXS simulations shown in blue [171] and orange [69], BAM simulations shown in green [68], and Einstein Toolkit [172] simulations shown in pink. Figure taken from Ref. [282].

by the observed behaviour of precessing NR where the coprecessing frame continues to precess roughly around the direction of final spin χ_f after merger [18], though this does not entail any systematic fitting to precessing NR. Therefore, incorporation of strong-field precession information into semi-analytic waveforms through calibration to precessing NR will be crucial to improve the accuracy of these models and reduce model systematics.

Let us therefore now consider analogously calibrating a precessing semi-analytic model in the same way as described above. As described in Sec. 2.4.2, the orbital precession dynamics is sourced by the four spin components instantaneously orthogonal to the orbital angular momentum (within the instantaneous orbital plane), and so now the waveform depends on all six spin degrees of freedom of the binary, as well as the mass ratio, which makes a total of seven intrinsic dimensions to calibrate to NR. If we assume similarly to the example of PhenomXHM above, that ten calibration points are needed for each of these dimensions, this results in requiring a seven-dimensional fit across 10^7 points in parameter space. This naïve approach to calibrating precessing semi-analytic waveforms is therefore

unfeasible for two reasons, the first being that precessing NR waveforms do not yet exist in numbers this large. Precessing NR simulations are computationally challenging to produce in large part due to their high dimensionality and lack of physical symmetries to reduce computational cost, and this is particularly true at high mass ratios. The largest publicly available catalogue of simulations to date is from the SXS collaboration [69], which contains 2018 simulations, of spin magnitudes up to $|\chi| = 0.998$ but mass ratios only up to $q \leq 10$. The second problem with calibrating to precessing NR comes from the high dimensional fitting procedure itself, as the hierarchical procedure currently used by PhenomXHM relies on ansätze which become incrementally more complicated as more intrinsic dimensions are included. This approach depends on a strategy of fitting first in the dimensions or parameter combinations which affect the waveform most strongly, and this is obfuscated by the high dimensionality of precessing binaries. Both these problems, therefore, motivate a dimensional reduction strategy for precessing binaries, in order to calibrate to semi-analytic models to NR.

The effective precession spin χ_p as introduced in Eq. 3.15 may present a natural way for calibrating precession effects in the strong-field through a single scalar parameter via the following effective mapping at some reference time t_{ref} :

$$\vec{\chi}_1(t_{\text{ref}}) = (\chi_{1x}, \chi_{1y}, \chi_{1z}) \mapsto \vec{\chi}'_1 = (\chi_p, 0, \chi_{1z}), \quad (4.1)$$

$$\vec{\chi}_2(t_{\text{ref}}) = (\chi_{2x}, \chi_{2y}, \chi_{2z}) \mapsto \vec{\chi}'_2 = (0, 0, \chi_{2z}), \quad (4.2)$$

where the spin components are defined in a Cartesian binary source frame with $\hat{L}(t_{\text{ref}}) = \hat{z}$. Such an identification reduces the four in-plane spin components to a single scalar quantity, making the problem of incorporating precession effects more tractable while capturing the dominant precession effects in the waveforms. This approach has successfully been implemented in the widely used phenomenological waveform approximant IMRPhenomPv2 [165], though in the current generation of precessing models, higher order PN spin information is included in the twisting up procedure to replicate fully precessing

waveforms, as described in Chapter 3. More recently, χ_p has been used in Ref. [272] for calibrating the non-precessing phenom model **PhenomD** [265, 266] to precessing NR. The dominant quadrupolar modes in the coprecessing frame and the Euler angles have been tuned to a relatively small set of precessing NR waveforms by reducing the precessing spin dimensions to a one-dimensional χ_p -based parameter which varies with mass ratio. We note that accuracy improvements as measured by mismatch for this model over its non-calibrated counterpart are modest, and it does not include calibration for higher modes.

The efficacy of such a χ_p -parameterisation has only been demonstrated in the inspiral [174] focusing on the (2, 2)-mode. Higher modes, however, are particularly important in binaries with large mass and spin asymmetries, for which also precession effects are more pronounced. While the radiation from a nonprecessing binary is dominated by the quadrupolar (2, 2)-mode, which is predominantly emitted along \hat{L} , in a precessing system power is transferred from the (2, 2)-mode to HMs. These HMs can become comparable in strength to the quadrupolar mode in the later inspiral and merger, and some modes, especially the (2, ± 1)-modes, can be particularly strong [173]. Additionally, a simple χ_p -parameterisation fails to accurately reproduce precession-induced mode mixing and the asymmetry between positive and negative m -modes [120]. In precessing systems where the relative power in HMs can be comparable to the dominant quadrupolar mode, this can lead to significant systematic errors [19, 48], as recent observations are starting to indicate [38, 210]. Therefore, the accurate modelling of HMs is particularly important in precessing systems. We will show in Sec. 4.3.1 that the simple χ_p -parameterisation of Eqs. (4.2) fails to accurately reproduce the behavior of precessing HMs, motivating the introduction of a new effective precession spin vector $\vec{\chi}_\perp$ which incorporates spin effects from both black holes, to address this issue. Focusing on the strong-field regime, we show that a $\vec{\chi}_\perp$ -parameterisation, (i) matches the opening angle of the precession cone at a given reference time; (ii) significantly better reproduces HMs than a χ_p -parameterisation; (iii) more accurately mimics the precession dynamics, and (iv) matches the final state. This

vectorial effective spin mapping could facilitate more accurate waveform modelling of precession in the strong-field regime.

4.1.1 A New Effective Precession Spin

To aid the calibration of the precessing sector of semi-analytic IMR waveform models, we seek to capture the dominant behavior through dimensional reduction by reducing the number of in-plane spin components through an effective map. To this end, we introduce a new dimensionless effective precession spin vector, $\vec{\chi}_\perp(t) \in \mathbb{R}^2$.

Our starting point for the construction of $\vec{\chi}_\perp$ is the opening angle of the precession cone at a reference time $t = t_{\text{ref}}$, $\lambda_L(t_{\text{ref}})$ given by Eq.(2.99), which captures the amount of precession in the system. We recall that the opening angle depends explicitly on the in-plane spin components through $S_\perp(t)$; we therefore seek a mapping such that this quantity is approximately preserved at the reference time at which the mapping is applied. To do so, we first place the in-plane spin projection of the total spin of the system onto the larger black hole, such that

$$\vec{\chi}_{1\perp}(t_{\text{ref}}) \mapsto \vec{S}_\perp(t_{\text{ref}})/m_1^2, \quad \vec{\chi}_{2\perp}(t_{\text{ref}}) \mapsto \vec{0}, \quad (4.3)$$

where

$$\vec{S}_\perp(t_{\text{ref}}) = m_1^2 \vec{\chi}_{1\perp}(t_{\text{ref}}) + m_2^2 \vec{\chi}_{2\perp}(t_{\text{ref}}). \quad (4.4)$$

We find, however, that this mapping can be further improved by assigning it conditionally to either the primary or secondary BH, depending on which BH has the largest in-plane spin magnitude $S_{i\perp}(t_{\text{ref}})$ at the reference time. This conditional placement ensures that a binary with an in-plane spin on only one BH is correctly reproduced. Furthermore, we impose the Kerr limit on the BH spin by including appropriate normalization factors into the definition of $\vec{\chi}_\perp$. With these constraints, we obtain the following effective precession

spin vector $\vec{\chi}_\perp(t_{\text{ref}})$:

$$\vec{\chi}_\perp(t_{\text{ref}}) \equiv \begin{cases} \frac{\vec{S}_\perp}{m_1^2 + S_{2\perp}}, & \text{if } S_{1\perp} \geq S_{2\perp}, \\ \frac{\vec{S}_\perp}{m_2^2 + S_{1\perp}}, & \text{if } S_{1\perp} < S_{2\perp}, \end{cases} \quad (4.5)$$

where the quantities $S_{1\perp}$, $S_{2\perp}$, \vec{S}_\perp are all evaluated at t_{ref} . We stress that the mass ratio q and the spin components along $\hat{L}(t_{\text{ref}})$ remain unaltered in this particular mapping. Explicitly, in a Cartesian binary source frame with $\hat{L}(t_{\text{ref}}) \equiv \hat{z}$, if $S_{1\perp}(t_{\text{ref}}) \geq S_{2\perp}(t_{\text{ref}})$ then we have

$$\vec{\chi}_1 = (\chi_{1x}, \chi_{1y}, \chi_{2z}) \mapsto \vec{\chi}_1' = (\chi_{\perp x}, \chi_{\perp y}, \chi_{1z}) \quad (4.6)$$

$$\vec{\chi}_2 = (\chi_{2x}, \chi_{2y}, \chi_{2z}) \mapsto \vec{\chi}_2' = (0, 0, \chi_{2z}), \quad (4.7)$$

where $\vec{\chi}_\perp = \{\chi_{\perp x}, \chi_{\perp y}\} = \{S_{\perp x}/(m_1^2 + S_{2\perp}), S_{\perp y}/(m_1^2 + S_{2\perp})\}$. Conversely, if $S_{1\perp}(t_{\text{ref}}) < S_{2\perp}(t_{\text{ref}})$ this mapping gives

$$\vec{\chi}_1 = (\chi_{1x}, \chi_{1y}, \chi_{2z}) \mapsto \vec{\chi}_1' = (0, 0, \chi_{1z}) \quad (4.8)$$

$$\vec{\chi}_2 = (\chi_{2x}, \chi_{2y}, \chi_{2z}) \mapsto \vec{\chi}_2' = (\chi_{\perp x}, \chi_{\perp y}, \chi_{2z}), \quad (4.9)$$

where $\vec{\chi}_\perp = \{\chi_{\perp x}, \chi_{\perp y}\} = \{S_{\perp x}/(m_2^2 + S_{1\perp}), S_{\perp y}/(m_2^2 + S_{1\perp})\}$. We note that instead of Cartesian coordinates, polar coordinates may be chosen. Then, $|\vec{\chi}_\perp|$ represents the magnitude of the mapped dimensionless spin vector and the azimuthal orientation, ϕ_\perp , is its angular position within the orbital plane at the reference time. We demonstrate the efficacy of this vectorial parameterisation, in particular for HMs, in Sec. 4.3.

To see the effect of this parameter reduction on the mode behaviour for a precessing binary, we present an example in Fig. 4.2. We show the $(2, \pm 1)$ -modes for a fiducial precessing binary with $q = 1.4$ and initial spins $\vec{\chi}_1 = (0.075, 0.043, 0.05)$, $\vec{\chi}_2 = (-0.346, 0.6, -0.4)$. The waveforms were produced using the NR surrogate waveform model `NRSur7dq4` [160] as




Figure 4.2: Amplitude of the $(2, 1)$ -mode (left) and the $(2, -1)$ -mode (right) for a fiducial precessing binary with $q = 1.4$, $\vec{\chi}_1(t_0) = (0.075, 0.043, 0.05)$, and $\vec{\chi}_2(t_0) = (-0.346, 0.6, -0.4)$. The fully precessing signal waveform is shown in blue, and the template waveforms parameterised by χ_p and $\vec{\chi}_\perp$ are shown in purple and orange, respectively. The $\vec{\chi}_\perp$ reproduces the phenomenology of this mode markedly better than the χ_p -mapping, especially in the merger-ringdown portion of the waveform.

further described in Sec. 4.2.1, with the orange modes produced using all six spin degrees of freedom, the modes shown in teal represent the same binary when the four in-plane spin components are replaced by a $\vec{\chi}_\perp$ -mapping as in Eqs. (4.6-4.9), and the modes in purple when the in-plane spin components are replaced with a χ_p -mapping as in Eq. (4.2). We find that $\vec{\chi}_\perp$ captures the fully precessing modes significantly better than a simple χ_p parameterisation. In particular, we see that unlike χ_p , $\vec{\chi}_\perp$ reproduces the amplitude and phasing of the mode oscillations on the orbital timescale and the amplitude modulations on the precession timescale much more faithfully. Additionally, amplitude modulations in the ringdown signal, which are completely missed in the χ_p -parameterisation, are much better captured. We note that the precession of this fiducial binary is dominated by the secondary BH spin – a region in the spin parameter space where χ_p knowingly performs poorly.

We next illustrate the impact of the conditional placement on this same fiducial binary of Figure 4.2, but this time showing the effect of an analogous conditional placement with the χ_p -parameterisation. As χ_p is known to perform poorly in this region of parameter




Figure 4.3: Amplitude of the (2,1)-mode (left), and time evolution of the two quaternion components, $q_1(t)$ and $q_2(t)$ (right), for the same fiducial binary as in Figs. 4.2 and 4.14, but with χ_p conditionally placed on the secondary BH. The fully precessing waveform mode and quaternions are shown in orange, with the system parameterized by a conditionally placed χ_p shown in purple, and $\vec{\chi}_\perp$ in teal. We see that although the conditional placement of χ_p does lead to an improvement in the accuracy with which it reproduces the (2,1)-mode, $\vec{\chi}_\perp$ still outperforms χ_p . Additionally, conditionally placing χ_p does not improve the accuracy with which it reproduces the precession dynamics, with $\vec{\chi}_\perp$ still much more closely matching the time evolution of the fully precessing \hat{q}_1 and \hat{q}_2 quaternion components.

space dominated by secondary spin, it is of interest to see whether an analogous conditional placement of the χ_p -mapping improves agreement of the modes for this binary. In Fig. 4.3, we show the fully precessing fiducial binary in orange, $\vec{\chi}_\perp$ -mapped system in teal, and the χ_p -mapped system with conditional placement in purple. We show the (2,1)-mode in the left panel, as well as two of the four quaternion component evolutions in the right panel. These quaternions are an alternative representation of the Euler angles to describe the precession dynamics, as described further in Subsec. 4.2.1. We note that now both $\vec{\chi}_\perp$ and χ_p are placed on the secondary BH for this binary. We can see in the left panel that the conditional placement does improve the accuracy with which χ_p reproduces the (2,1)-mode; however the $\vec{\chi}_\perp$ -mapping still reproduces the phasing of the mode significantly better. Additionally, we see in the right panel that the conditional placement of χ_p does not produce the precession dynamics more accurately, with $\vec{\chi}_\perp$ still much more closely matching the time evolution of the fully precessing quaternion components.

We present two additional examples of waveform modes and corresponding precession dynamics produced using the effective spin mappings, both χ_p and $\vec{\chi}_\perp$. For each of the binaries considered here, we show the (2,1)-mode, as well as two of the four quaternion component evolutions, as above in Fig. 4.3. In all of the figures for these examples, the fully precessing system's results are shown in orange, the results parameterized by $\vec{\chi}_\perp$ in teal, and by χ_p in purple.




Figure 4.4: Amplitude of the (2,1)-mode (left), and time evolution of the two quaternion components, $q_1(t)$ and $q_2(t)$ (right), for an equal-mass binary with initial spins $\chi_1(t_0) = (0.225, 0.13, -0.15)$, and $\chi_2(t_0) = (0.09, 0.15, 0.1)$. We show the fully spinning waveform mode and quaternion components in orange. The mode and quaternions parameterized by χ_p are shown in purple, and $\vec{\chi}_\perp$ in teal. We see that $\vec{\chi}_\perp$ more faithfully reproduces the fully precessing (2,1)-mode, and much more accurately reproduces the precession dynamics, than χ_p .

The next binary we consider is equal-mass, i.e. $q = 1$, with initial spins $\vec{\chi}_1(t_0) = (0.225, 0.13, -0.15)$, $\vec{\chi}_2(t_0) = (0.09, 0.15, 0.1)$. In this equal-mass limit, we expect χ_p to perform poorly, and $\vec{\chi}_\perp$ to perform much better, as discussed more thoroughly in Sec. 4.4. Indeed, in the left panel of Figure 4.4, we see that $\vec{\chi}_\perp$ better replicates the (2,1)-mode for this particular binary, with an amplitude closer to that of the fully precessing waveform and slightly improved phasing. We see the improvement by using $\vec{\chi}_\perp$ as opposed to χ_p more clearly in the dynamics as shown in the right panel. The good agreement between the time evolution of fully precessing quaternion components in orange and those of the $\vec{\chi}_\perp$ -mapped system in teal, is in stark contrast to the χ_p -mapped components in purple,

which matches the dynamics poorly.




Figure 4.5: Amplitude of the $(2,1)$ -mode (left), and time evolution of the two quaternion components, $q_1(t)$ and $q_2(t)$ (right), for a $q = 3$ fiducial binary, with initial spins $\vec{\chi}_1(t_0) = (0.45, 0.26, -0.3)$ and $\vec{\chi}_2(t_0) = (0.15, 0.087, 0.1)$. The fully precessing waveform mode and quaternions are shown in orange, with the system parameterized by χ_p shown in purple, and $\vec{\chi}_\perp$ in teal. We see that $\vec{\chi}_\perp$ more faithfully reproduces the fully precessing $(2,1)$ -mode, capturing the correct phasing of the mode, and more accurately reproduces the precession dynamics represented by the quaternion components, than χ_p .

Our final example is a binary with $q = 3$, and initial spins $\vec{\chi}_1(t_0) = (0.45, 0.26, -0.3)$, $\vec{\chi}_2(t_0) = (0.15, 0.08, 0.1)$. We note that in this example, unlike the fiducial binary shown in Figure 4.2, $\vec{\chi}_\perp$ is mapped onto the primary BH. The left panel of Figure 4.5 shows the $(2,1)$ -mode, where we see that unlike in the previous example, the amplitudes of the two mapped waveform modes are very similar to that of the fully precessing mode (orange). However, $\vec{\chi}_\perp$ clearly much better matches the phasing of the fully precessing mode, with the orange and teal lines being indistinguishable for much of the inspiral, in contrast to χ_p which shows a clear dephasing, especially in the merger ringdown. We also see that, like for the other fiducial binaries, the quaternion components mapped by $\vec{\chi}_\perp$, much more faithfully represent the fully precessing quaternions, compared to the χ_p -mapped components, demonstrating that $\vec{\chi}_\perp$ better replicates the precession dynamics of the fully precessing system. Therefore although the precession dominated by the primary BH spin in Figure 4.5, where we still observe noticeable improvements with $\vec{\chi}_\perp$ over χ_p .

4.2 Methodology

4.2.1 Waveforms

To systematically assess the efficacy of the new effective parameterisation Eq. (4.5), we compare the waveforms obtained from the seven-dimensional system characterized by $(q, \vec{\chi}_1, \vec{\chi}_2)$ to the five-dimensional effective system described by $(q, \chi_{1\parallel}, \chi_{2\parallel}, \vec{\chi}_\perp)$ as well as to the four-dimensional system given by $(q, \chi_{1\parallel}, \chi_{2\parallel}, \chi_p)$. As we are particularly interested in testing the efficacy of such mappings in the strong-field regime, we use the NR surrogate model `NRSur7dq4` [160] as provided through the public `GWSURROGATE` PYTHON package [175] to generate late inspiral-merger-ringdown waveforms for all our analyses. This is a surrogate model built over a set of existing NR simulations as described in Sec. 3.5, to interpolate over them in a pre-defined parameter space. As it is based upon NR, it therefore contains the full precession and spin dynamics as well as their imprint on the waveforms, with no spin dimension reductions or effective parameterisations. The computational efficiency of this model allows us to assess the mappings over a dense sampling of the intrinsic parameter space. However, due to the limited parameter ranges of the NR simulations it is built upon, the surrogate is limited to dimensionless spin magnitudes $|\chi_i| \leq 0.8$ and mass ratios $q \leq 4$. While precession effects are even more pronounced for higher mass ratios, the importance of the in-plane spin on the smaller BH decreases and therefore we expect any dimensional reduction that is built to capture the dominant precession spin to perform even better in this limit.

The surrogate model represents an interpolant across a discrete set of NR simulations [89, 90, 152]. The precessing waveform modes up to $\ell \leq 4$ are obtained by following the strategy outlined in Sec. 3.1, where the coprecessing modes are further decomposed into coorbital modes to further simplify their structure,

$$h_{\ell m}^{\text{coprec}}(t) = e^{im\Omega(t)} h_{\ell m}^{\text{coorb}}(t), \quad (4.10)$$

where $\Omega(t)$ is the relative angular velocity relating the two frames [91].

Unlike most other waveform models, the surrogate makes use of four unit quaternion components $\{\hat{q}_0(t), \hat{q}_1(t), \hat{q}_2(t), \hat{q}_3(t)\}$ instead of three Euler angles to describe the precession dynamics of the orbital plane [125]. Importantly, the precessing modes are obtained in an inertial frame corresponding to $\hat{L}(t_0) \equiv \hat{z}$ at the initial time t_0 , as opposed to the more commonly used \hat{J} -aligned frame. In this coordinate frame, the xy -plane coincides with the initial orbital plane of the binary with the x -axis parallel to the separation vector pointing from the smaller black hole to the larger one. Due to this binary source frame choice, caution must be taken when interpreting the physical meaning of the quaternions, as they do not represent a transformation to the commonly-used frame ‘J-frame’, where the initial z -axis is aligned with the total angular momentum.

The resulting waveforms are of a fixed length, from $t_0 = -4300M$, the negative sign indicating premerger, up to $t = +100M$ after the merger. This relatively short length makes them unsuitable for describing the waveforms of low mass binaries $M \lesssim 70M_\odot$ assuming a starting frequency of 20Hz. The surrogate determines the coalescence time t_c as the peak of the quadrature sum of the mode amplitudes,

$$t_c = \max_t \sqrt{\sum_{\ell m} |h_{\ell m}(t)|^2}, \quad (4.11)$$

and shifts the time arrays such that the peak amplitude occurs at $t_c = 0$.

4.2.2 Faithfulness for Precessing Waveforms

For our quantitative comparisons, we define h to represent the fully precessing waveform with all 6 spin degrees of freedom, and h_σ the corresponding waveform produced by an effective mapping, where $\sigma \in [\chi_p, \vec{\chi}_\perp]$. Hereafter, we refer to h as the signal waveform, and to h_σ as the template waveform. The effective mappings are applied at the surrogate initial time t_0 , such that the full and mapped spins are used as initial data to produce the signal and template waveforms respectively. To quantify how well either mapping

reproduces the full waveform, we compute the match (faithfulness) between h and h_σ . We defined the match in Eq. (2.74) to the noise-weighted inner product between two waveforms maximized over a time and phase shift of the template waveform. To be concrete, in this paper we will compute matches between the signal waveform h and template waveform h_σ :

$$\mathcal{M}(h, h_\sigma) = \max_{t_{c\sigma}, \phi_{0\sigma}} \frac{\langle h, h_\sigma \rangle}{\sqrt{\langle h, h \rangle \langle h_\sigma, h_\sigma \rangle}}, \quad (4.12)$$

where the inner product is defined as in Eq. (2.75). In what follows, h and h_σ either denote individual waveform modes $h_{\ell m}$ defined as in Eq. (3.3), or the complex strain.

To assess how accurately individual modes, in particular HMs, are reproduced under the effective mapping, we compute individual mode-by-mode matches between each spin mapping and the full-spin waveform; i.e. for each pair (ℓ, m) , h , h_σ in Eq. (4.12) are replaced by individual modes $h \rightarrow h_{\ell m}$ and $h_\sigma \rightarrow h_{\sigma, \ell m}$.

As the odd m -modes are sourced by mass and spin asymmetries, they are often contaminated by numerical noise for systems with small asymmetries. We therefore employ an additional cut on the energy $E_{\ell m}$ contained in the inertial-frame $(2, \pm 1)$ - and $(3, \pm 3)$ -modes of the fully spinning mode prior to calculating the match, where the mode energy is given by

$$E_{\ell m} = \frac{1}{16\pi} \int_{t_0}^{t_f} |\dot{h}_{\ell m}(\tau)|^2 d\tau, \quad (4.13)$$

where t_f is the final time of the surrogate waveforms. Based on calculations of the energy contained in those modes for binaries without mass or spin asymmetries, we find the energy thresholds for these modes given by the values listed in Table 4.1. Modes with $E_{\ell m}$ less than these values are discarded in the mode-by-mode match calculations performed in Sec. 4.3.1.

We perform mode-by-mode match calculations using both white noise, i.e. $S_n(f) = 1$ and the projected aLIGO PSD for the fourth observing run [51], denoted $\mathcal{M}_{\text{white}}$ and \mathcal{M}_{O4} respectively. The white noise matches are to assess the systematic errors induced by

(ℓ, m) -mode	$E_{\ell m}$ threshold
$(2, \pm 1)$	1.0×10^{-6}
$(3, \pm 3)$	5.5×10^{-7}

Table 4.1: Mode energy thresholds for odd m -modes. If the energy of a particular mode is below its threshold, the mode is considered to be numerical noise and excluded from the mode-by-mode match calculation.

the mappings in the absence of detector-specific frequency sensitivities, while the PSD-weighted matches demonstrate the effect for a given detector. For the detector PSD matches we choose a starting frequency of $f_{\min} = 20\text{Hz}$, and truncate the waveforms at $t = 50M$ after the peak as determined by Eq. (4.11) to remove postmerger numerical noise.

While the individual mode matches allow us to assess how well HMs in particular are captured by the lower-dimensional spin parameterisation, GW detectors measure the strain, which also depends on extrinsic parameters of the source such as the luminosity distance, the effective polarization angle κ [176] and the binary inclination ι relative to the line of sight of an observer.

Following Refs. [19, 48, 126], we compute the strain match by analytically optimizing over the template polarization angle κ_σ and numerically optimizing over the template reference phase $\phi_{0\sigma}$ and template coalescence time $t_{c\sigma}$,

$$\mathcal{M}_{\text{strain}}(M, \iota, \phi_0, \kappa) = \max_{t_{c\sigma}, \phi_{0\sigma}, \kappa_\sigma} \left. \frac{\langle h, h_\sigma \rangle}{\sqrt{\langle h, h \rangle \langle h_\sigma, h_\sigma \rangle}} \right|_{\iota=\iota_\sigma}. \quad (4.14)$$

We do not optimize over any intrinsic parameters. We note that Eq. (4.14) still depends on the signal polarization κ and reference phase ϕ_0 . By averaging over these two angles, we obtain the sky-and-polarization-averaged strain match,

$$\overline{\mathcal{M}}_{\text{strain}}(M, \iota) = \frac{1}{8\pi^2} \int_0^{2\pi} d\kappa \int_0^{2\pi} d\phi_0 \mathcal{M}_{\text{strain}}(M, \iota, \phi_0, \kappa). \quad (4.15)$$

Additionally, to account for the correlation between low matches and low signal-to-

noise ratio (SNR), we also compute the SNR-weighted strain match [19, 48] given by

$$\overline{\mathcal{M}}_{\text{SNR}}(M, \iota) = \left(\frac{\sum_i (\mathcal{M}(h, h_\sigma))^3 \langle h_i, h_i \rangle^{3/2}}{\sum_i \langle h_i, h_i \rangle^{3/2}} \right)^{1/3}, \quad (4.16)$$

where the sum is over a discrete range of source polarizations κ and initial phases ϕ_0 as detailed in Sec. 4.2.3.

We note that we do not apply the postmerger truncation at $t = 50M$, nor do we impose the mode energy thresholds of Table 4.1 when computing the sky-and-polarization-averaged and the SNR-weighted strain matches. For strain matches we take into account all modes up to $\ell = 4$ as provided by the NR surrogate.

Lastly, rather than showing the agreement between two waveforms, it can be advantageous to quantify the disagreement through the mismatch $\mathcal{M}\mathcal{M}$ instead:

$$\overline{\mathcal{M}\mathcal{M}}_{\text{strain}} = 1 - \overline{\mathcal{M}}_{\text{strain}}, \quad (4.17)$$

$$\overline{\mathcal{M}\mathcal{M}}_{\text{SNR}} = 1 - \overline{\mathcal{M}}_{\text{SNR}}. \quad (4.18)$$

4.2.3 Binary Configurations

The mode-by-mode matches are computed for a large number of mass ratios and spins that systematically sample the validity range of the surrogate model with the details provided in the second column of Table 4.2. We choose the initial time as the reference time, i.e. $t_{\text{ref}} \equiv t_0$, and sample the initial spins in a spherical coordinate system using the spin magnitudes $|\vec{\chi}_i|$, the azimuthal orientations $\phi_i = \arccos(\hat{S}_i \cdot \hat{x})$, and the cosine of the tilt angles $\cos(\theta_i) = \hat{S}_i \cdot \hat{L}$. Specifically, we keep the initial azimuthal orientation of the spin of the larger BH ϕ_1 of $\vec{\chi}_{1\perp}$ fixed, while rotating $\vec{\chi}_{2\perp}$ to achieve a range of angular azimuthal separations, and vary the initial tilt angles θ_i systematically. Further, we only choose configurations with at least one spinning BH, demanding that at least one BH has a nonzero in-plane spin, thereby excluding aligned-spin or nonspinning binaries. This amounts to a total of 47,136 unique binary configurations in terms of their intrinsic

4. A NEW EFFECTIVE PRECESSION SPIN

parameters $\{q, \vec{\chi}_1, \vec{\chi}_2\}$. When considering a detector PSD, we additionally consider three values of the total mass, 75, 150 and $250M_\odot$.

	$\mathcal{M}\mathcal{M}_{\ell m}$	$\mathcal{M}\mathcal{M}_{\text{strain}}$	$\mathcal{M}\mathcal{M}_{\text{SNR}}$
$\ \vec{\chi}_i\ $	$\ \vec{\chi}_i\ \in [0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8]$, if $\ \vec{\chi}_1\ = 0$, $\ \vec{\chi}_2\ \neq 0$	$\ \vec{\chi}_i\ \in U[0, 0.8]$	$\ \vec{\chi}_1\ , \ \vec{\chi}_2\ \in U[0, 0.8]$
θ_i [rad]	$\theta_i \in [0, \pi/3, 2\pi/3, \pi]$, if $\theta_1 \in [0, \pi]$, $\theta_2 \notin [0, \pi]$	$\cos(\theta_i) \in U[-1, 1]$	$\cos(\theta_i) \in U[-1, 1]$
ϕ_i [rad]	$\phi_1 = \pi/3$, $\phi_2 \in [0, \pi/3, 2\pi/3, \pi, 4\pi/3, 5\pi/3]$	$\phi_i \in U[0, 2\pi]$	$\phi_i \in U[0, 2\pi]$
q	$q \in [1, 1.5, 2, 3, 4]$	$q \in U[1, 4]$	$q \in U[1, 4]$
M_{tot} [M_{\odot}]	$M_{\text{tot}} \in [75, 150, 250]$ (PSD matches only)	$M_{\text{tot}} \in U[70, 250]$	$M_{\text{tot}} \in [75, 100, 125, 150, 175, 200, 225, 250]$
ι [rad]	–	$\cos(\iota) \in U[-1, 1]$	$\iota = \pi/3$
ϕ_0 [rad]	–	$\phi_0 \in U[0, 2\pi]$	$\phi_0 \in [0, \pi/4, \pi/2, 3\pi/4, \pi, 5\pi/4, 3\pi/2, 7\pi/4]$ (included in sum)
κ [rad]	–	$\kappa \in U[0, \pi/4]$	$\kappa \in [0, \pi/12, \pi/6, \pi/4]$ (included in sum)
Total binaries sampled	$\mathcal{M}\mathcal{M}_{\ell m, \text{white}}$: 47,136 $\mathcal{M}\mathcal{M}_{\ell m, \text{PSD}}$: 141,408	20,833	...

Table 4.2: Binary configurations used in the different match calculations. Binaries for mode-by-mode match calculations are sampled systematically across the intrinsic parameter space, and three total mass scales are used for the O4 PSD matches. Since it is the relative azimuthal separation of spins which is important, we choose to keep $\phi_1 = \pi/3$ fixed while changing ϕ_2 . Additionally, we place constraints on the spin magnitudes and tilt angles such that none of our binary configurations have both BHs with aligned spins or non-spinning. For the strain matches, the intrinsic parameters are drawn from random uniform distributions (shown by $U[a, b]$ in the table); the SNR-weighted matches use fixed extrinsic parameter values while they are drawn randomly for the sky-and-polarization-averaged matches.

For the strain matches as given in Eqs. (4.15) and (4.16), additional extrinsic parameters, namely binary inclination ι , initial phase ϕ_0 , and polarization κ , need to be taken into account. In such high dimensions, systematic sampling becomes unfeasible. Therefore, for the sky-and-polarization-averaged matches we draw the intrinsic and extrinsic binary parameters from random uniform distributions as detailed in the third column of Table 4.2, considering a total of 20,833 unique binary configurations.

For the SNR-weighted strain matches, we first draw 100 binary configurations randomly from the 20,833 used to compute the sky-and-polarization-averaged strain matches,

only considering their intrinsic parameters, $\{q, |\vec{\chi}_1|, |\vec{\chi}_2|, \cos(\theta_1), \cos(\theta_2), \phi_1, \phi_2\}$. We fix the source inclination at a moderate inclination of $\iota = \pi/3$. As detailed in the last column of Table 4.2, for each binary configuration we choose eight initial phase and four polarization values and compute 32 matches $\mathcal{M}(h, h_\sigma)$, one for each pair $\{\phi_0, \kappa\}$, which are then summed into a single SNR-weighted match for each binary configuration as per Eq. (4.16). We repeat this calculation for each of the total masses detailed in Table 4.2, noting that we use the same 100 intrinsic binary configurations for each M . This yields 800 SNR-weighted strain matches for each mapping $\sigma \in \{\chi_p, \vec{\chi}_\perp\}$.

4.3 Results

4.3.1 Mode Analysis

We first assess how well the vectorial effective spin parameter $\vec{\chi}_\perp$ reproduces individual modes, in particular HMs, for different mass ratios. To quantify the efficacy of $\vec{\chi}_\perp$ across the parameter space, we compute white noise mode-by-mode matches for the $(2, \pm 2)$ -modes and a selection of HMs for all binaries listed in the second column of Table 4.2. Figure 4.6 shows the cumulative match results for the $(2, 2)$ - and $(2, 1)$ -modes for mass ratio $q = 1$ and $q = 3$.

We expect the χ_p -parameterisation to perform well at replicating the dominant $(2, \pm 2)$ -mode behavior, and indeed we see similar results in this mode for both parameterisations, if slightly improved with the new effective spin, except for the equal-mass case, where we find a more marked improvement. We attribute this to the fact that χ_p is designed to replicate the average precession rate, but in equal-mass configurations the in-plane spin vectors precess at the same rate and become orientationally locked, which is not captured correctly by χ_p [132, 174]. Additionally, $\vec{\chi}_\perp$ takes into account the in-plane spins on both black holes, while χ_p selects only the larger spin component leading to a systematic underestimation of the total in-plane spin for equal-mass cases.

We observe the most dramatic improvements in the $(2, \pm 1)$ -modes. For example, for

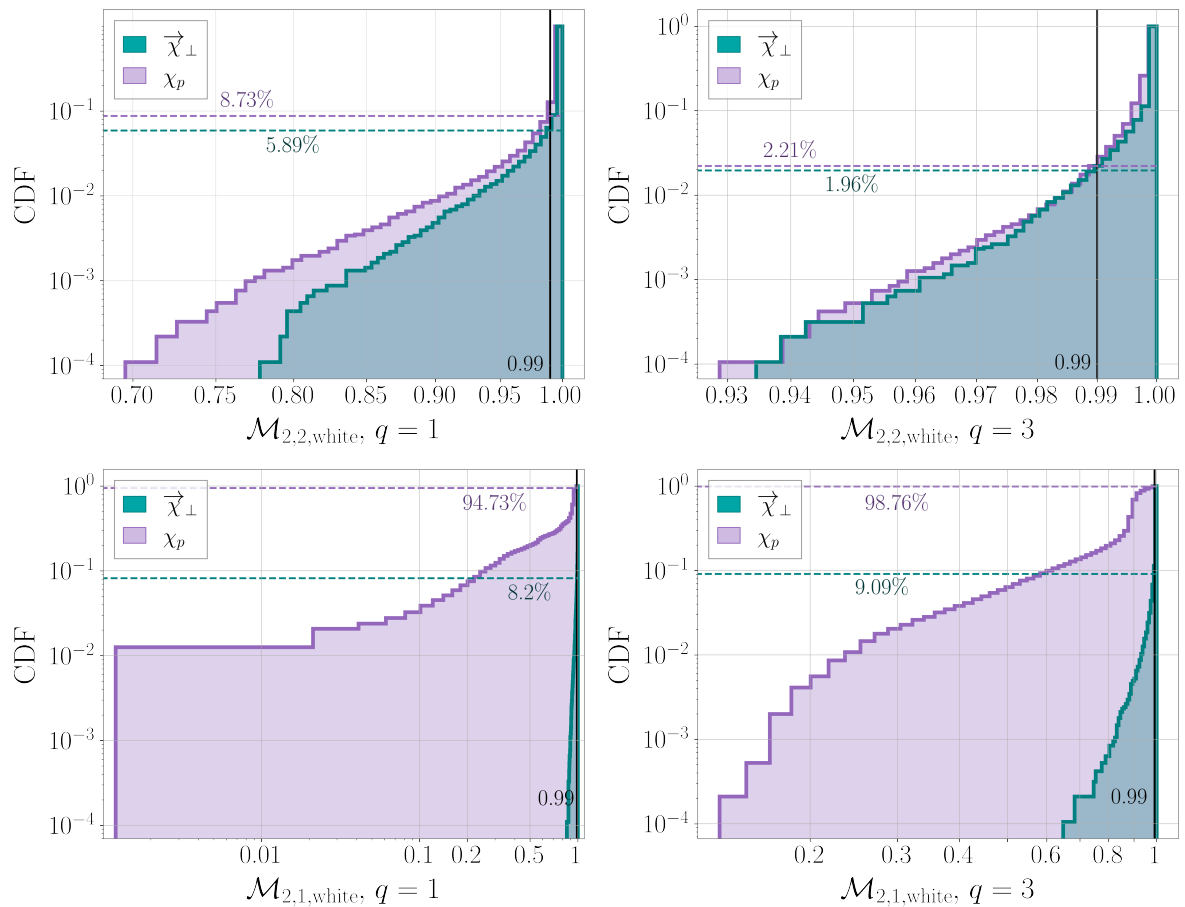


Figure 4.6: Cumulative histograms of white noise matches for the (2,2)-mode (top row) and the (2,1)-mode (bottom row) for 9,120 binaries with mass ratios $q = 1$ (left column) and 9,504 binaries with $q = 3$ (right column). Details of how these binaries are systematically sampled can be found in Table 4.2. Teal histograms show the results using the $\vec{\chi}_\perp$ -parameterisation, purple ones χ_p . The dashed horizontal lines indicate the percentage of matches below 0.99. Using $\vec{\chi}_\perp$, we see a clear improvement over the χ_p -mapping for the (2,1)-mode, and comparable if slightly better performance for the (2,2)-mode. Results for additional modes and mass ratios are presented in Figure 4.7.

$q = 3$ shown in Figure 4.6, the percentage of matches below 0.99 decreases dramatically from 94.73% with χ_p to 8.2% with $\vec{\chi}_\perp$. Note that the long tails toward very low matches for the χ_p -parameterisation, and the comparatively short ones of $\vec{\chi}_\perp$, are a generic feature across all HMs we analyzed, suggesting that $\vec{\chi}_\perp$ better replicates the higher mode behavior even when it performs at its worst.

Results of white noises matches for additional modes and mass ratios are shown in




Figure 4.7: Complete results for white noise matches between the fully spinning waveform and each of the effective spin parameterisations, $\vec{\chi}_\perp$ (teal) and χ_p (purple). We show percentages of matches split by mass ratio and mode, which fall below a threshold of 0.99. The dashed horizontal lines indicate the 1%, 10% and 50% marks. We see improved results using $\vec{\chi}_\perp$ as opposed to χ_p across all mass ratios and modes, and note particularly the dramatic improvements in performance for HMs, especially the $(2, \pm 1)$ -modes.

Fig. 4.7. We split the results by mass ratio $q \in [1, 1.5, 2, 3, 4]$ and mode $(\ell, m) \in [(2 \pm 2), (2 \pm 1), (3, \pm 3), (4, \pm 4)]$. We show the percentages of matches below 0.99 between the fully spinning waveform, and the waveforms produced by each of the two effective spins, $\vec{\chi}_\perp$ in teal, and χ_p in purple. We first note that at the match threshold of 0.99, we see an improvement by using $\vec{\chi}_\perp$ over χ_p , across all mass ratios and modes. These improvements are particularly dramatic for higher modes, particularly the $(2, \pm 1)$ -modes. For example, at mass ratio $q = 4$, the percentage of $(2, 1)$ -mode matches below 0.99 using the χ_p parameterisation is 96.3%, which improves dramatically with $\vec{\chi}_\perp$ to just 4.3%. The parameterisations perform more similarly for the $(2, \pm 2)$ -modes, but even for the quadrupolar modes we see small improvements. For example, at mass ratio $q = 1.5$ we see the percentage of $(2, 2)$ -mode matches below 0.99 improving from 8.5% with χ_p to 4.7% with $\vec{\chi}_\perp$.

Intriguingly, for the $(2, \pm 1)$ and $(4, \pm 4)$ -modes, both parameterisations perform worst at $q \sim 1.5$, after which their performance improves with increasing mass ratio. To further investigate this intermediate region between the equal-mass regime and higher mass ratios where the secondary spin becomes less important, we performed additional white noise matches at mass ratios $q \in [1.2, 1.4, 1.6, 1.8]$. We find that the performance of both spin mappings improves with increasing mass ratio for the $(2, \pm 2)$ - and $(3, \pm 3)$ -modes, with $\vec{\chi}_\perp$ consistently outperforming χ_p . We also find that $\vec{\chi}_\perp$ performs worst around $q \sim 1.4$ for the $(2, \pm 1)$ - and $(4, \pm 4)$ -modes, but that the distributions for both mappings are fairly flat between $q = 1.2$ and $q = 2$, and even at its worst $\vec{\chi}_\perp$ still vastly outperforms χ_p . For example, in the $(2, 1)$ -mode, at $q = 1.4$, the percentage of matches below 0.99 is 100% with χ_p , and only 26.35% with $\vec{\chi}_\perp$. We also note that we find only minor differences between the positive and negative m -modes for both mappings, and neither performs consistently worse at replicating either positive or negative m -modes.

Additionally, our new mapping shows moderate improvements for the $(3, \pm 3)$ -modes and striking improvements for the $(4, \pm 4)$ -modes, with the improvements particularly marked at equal-mass and at our highest mass ratio $q = 4$. In summary, we find that the $\vec{\chi}_\perp$ -parameterisation performs consistently better than χ_p for every mass ratio and across all modes, and in particular for odd m -modes.

In addition to the white noise matches, we repeat the analysis using the projected O4 aLIGO PSD [51] with $f_{\min} = 20\text{Hz}$. These matches also require a total mass scale, so we choose three masses $M \in \{75, 150, 250\}M_{\odot}$ compatible with the fixed length of the surrogate waveforms, bringing the total sampled binaries to 141,408. Full details of the sampling are given in Table 4.2. We obtain qualitatively similar results to the white noise matches, as shown in Figure 4.8. All of the matches improve slightly compared to the white noise matches across both mappings due to the frequency weighting of the PSD, but the features of our results and conclusions remain the same: The $\vec{\chi}_{\perp}$ -mapping significantly improves upon χ_p for the $(2, \pm 1)$ - and $(4, \pm 4)$ -modes, with moderate improvements for the $(3, \pm 3)$ -modes, and comparable if slightly better performance for the $(2, \pm 2)$ -modes.

4.3.2 Strain Analysis

In the previous section we have demonstrated the improvement of $\vec{\chi}_{\perp}$ over χ_p at the level of individual $h_{\ell m}$ -modes. We now assess the degree to which the improvement in the HMs impacts the strain. Figure 4.9 shows the strain for the fiducial binary at an inclination of $\iota = \pi/3$. The excellent agreement between the fully precessing waveform (orange) and the one parameterised by $\vec{\chi}_{\perp}$ (teal) throughout the late inspiral as well as the merger ringdown is clearly visible. To quantify this agreement, we first compute the sky-and-polarization-averaged strain mismatches for 20,833 binary configurations as detailed in Table 4.2 using the O4 PSD and $f_{\min} = 20\text{Hz}$. Our results for both effective parameterisations are shown in Figure 4.10. Using $\vec{\chi}_{\perp}$ rather than χ_p , we find a median improvement of more than 1 order of magnitude from 4×10^{-3} to 2×10^{-4} . Furthermore, we note the non-negligible tail of extremely low mismatches below 10^{-6} for $\vec{\chi}_{\perp}$.

As low matches are often correlated with low SNRs and, therefore, with a lower detection probability, we also compute the SNR-weighted mismatch Eq. (4.18) for 100 randomly drawn intrinsic binary configurations as given in the fourth column of Table 4.2 for a moderate inclination of $\iota = \pi/3$ at t_0 . Similar to the sky-and-polarization-averaged strain mismatches, we see an improvement of around 1 order of magnitude when using




Figure 4.8: Full results for the O4 PSD-weighted matches between the fully spinning waveform and each of the effective spin parameterisations, $\vec{\chi}_\perp$ (teal) and χ_p (purple), split by total mass, mass ratio, and mode, which have a match less than 0.99. The dashed horizontal lines indicate 1%, 10% and 50%. Similar to the white noise matches of Figure 4.7, we see improvements by using $\vec{\chi}_\perp$ over χ_p across all masses, mass ratios and modes. We note the dramatic improvements in performance for HMs, particularly the $(2, \pm 1)$ -modes, and also the high degree of similarity between each of the three total masses.




Figure 4.9: Amplitude of the waveform strain $h(t)$ for the same fiducial binary as in Figure 4.2 at an inclination of $\iota = \pi/3$. The figure shows the fully precessing waveform (orange) along with both the χ_p (purple) and $\vec{\chi}_\perp$ -parameterisations (teal). The $\vec{\chi}_\perp$ -mapping reproduces the strain amplitude much more faithfully than χ_p , especially in the late inspiral portion of the waveform. Note that in the late inspiral, where the blue line cannot be seen, it is indistinguishable from the orange line.

the $\vec{\chi}_\perp$ -parameterisation instead of χ_p , as shown in Figure 4.11. The worst two cases for each parameterisation are highlighted in both panels. We see that the worst cases for $\vec{\chi}_\perp$ (teal and blue) perform similarly under both mappings, if slightly better with the new $\vec{\chi}_\perp$ -mapping. These cases both have a mass ratio of $q \sim 1.5$, which as noted previously in Sec. 4.3.1, is a mass ratio where both parameterisations perform worst. The worst cases for χ_p (purple and orange) on the other hand, show significant improvements of around 2 and 4 orders of magnitude respectively across the entire mass range when the $\vec{\chi}_\perp$ -mapping is used.

To better understand these marked improvements we employ several diagnostics. First, we investigate whether there exists a correlation between the initial opening angle of the precession cone $\lambda_L(t_0)$ [Eq. (2.99)] and the sky-and-polarization-averaged strain mismatch. We define the difference in the initial precession cone opening angle between the mapped

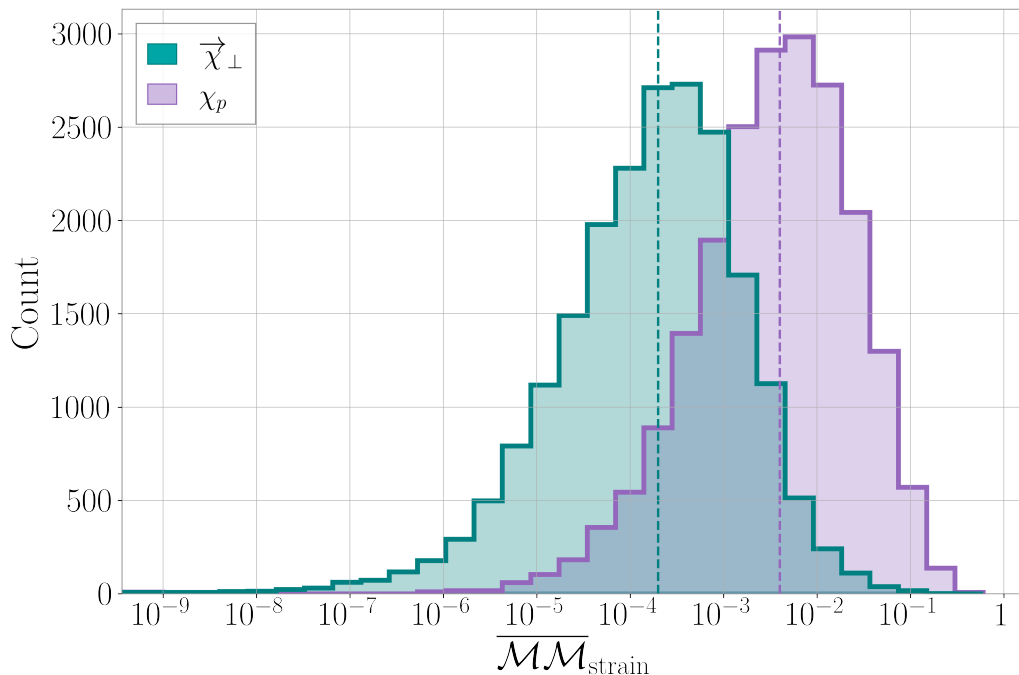


Figure 4.10: Histograms of the sky-and-polarization-averaged strain mismatches $\overline{\mathcal{M}\mathcal{M}}_{\text{strain}}$ between the fully precessing waveform and each of the two-spin mappings using the O4 PSD. The parameters for each of the 20,833 binaries tested were drawn from random uniform distributions as outlined in Table 4.2. The dashed vertical lines show the median mismatch for each mapping. We see an improvement in the median mismatch of more than one order of magnitude when using $\vec{\chi}_{\perp}$.

and unmapped system, $\Delta\lambda_L(t_0)$, as

$$\Delta\lambda_L(t_0) \equiv \lambda_L(t_0) - \lambda_{\sigma L}(t_0), \quad (4.19)$$

where λ_L is given by Eq. (2.99). The definition of $\lambda_{\sigma L}$ is the same as for λ_L , but replaces S_{\perp} with $S_{\sigma\perp}$, where S_{\perp} is the initial total in-plane spin magnitude before the mapping and $S_{\sigma\perp}$ is the total in-plane spin magnitude after the mapping. All quantities are evaluated at the initial time t_0 , and we approximate L by its Newtonian value $L = \mu\sqrt{Mr}$, where $\mu = m_1m_2/M$ is the reduced mass and $r = M^{1/3}\omega_{\text{orb}}^{-2/3}$ with ω_{orb} the orbital angular frequency.

Figure 4.12 shows $\Delta\lambda_L(t_0)$ against the strain mismatch for each of the 20,833 binaries, calculated with both the $\vec{\chi}_{\perp}$ (teal) and χ_p (purple) effective spin mappings. For the new

4. A NEW EFFECTIVE PRECESSION SPIN




Figure 4.11: SNR-weighted strain mismatches $\overline{\mathcal{M}\mathcal{M}}_{\text{SNR}}$ as a function of binary total mass M for 100 binaries for the $\vec{\chi}_\perp$ (left) and χ_p (right) mappings. The teal and blue lines correspond to the cases $q = 1.3$, $\vec{\chi}_1(t_0) = (0.37, -0.36, 0.46)$, $\vec{\chi}_2(t_0) = (-0.32, 0.02, 0.13)$ and $q = 1.6$, $\vec{\chi}_1(t_0) = (-0.21, 0.23, 0.58)$, $\vec{\chi}_2(t_0) = (-0.26, -0.56, 0.34)$ respectively, which show the worst results for the $\vec{\chi}_\perp$ -mapping. The purple and orange lines show the two cases $q = 3.2$, $\vec{\chi}_1(t_0) = (-0.66, 0.12, 0.01)$, $\vec{\chi}_2(t_0) = (0.23, -0.25, -0.03)$ and $q = 3.2$, $\vec{\chi}_1(t_0) = (-0.55, 0.09, 0.07)$, $\vec{\chi}_2(t_0) = (0.03, 0.04, 0.03)$ where χ_p shows the worst performance. We see an average improvement of around 1 order of magnitude using the $\vec{\chi}_\perp$ -mapping.

mapping, we see a clear correlation between lower values of $\Delta\lambda_L(t_0)$ and lower strain mismatches. Overall, the $\vec{\chi}_\perp$ -parameterisation yields a more accurate initial cone opening angle resulting in a more faithful representation of the fully precessing waveform.

The better agreement between the initial opening angles suggests that the spins themselves are captured more faithfully. To show this, as a second diagnostic we compare the spin evolutions of a fully precessing binary with its effective counterparts. Figure 4.13 shows the time evolution of the total in-plane spin S_\perp (orange) for the fully precessing fiducial binary and those of the $\vec{\chi}_\perp$ (teal) and χ_p (purple) parameterisations for the fiducial binary. We obtain these by transforming the spin evolutions in the inertial frame to the coprecessing frame using the quaternions. It is evident that the two-dimensional




Figure 4.12: Sky-and-polarization-averaged strain mismatch $\overline{\mathcal{M}\mathcal{M}}_{\text{strain}}$ versus $\Delta\lambda_L(t_0)$ for the 20,833 binaries with both the $\vec{\chi}_\perp$ (teal) and χ_p (purple) effective spin parameterisations. We observe a slight correlation between small $\Delta\lambda_L(t_0)$ and low mismatches for $\vec{\chi}_\perp$, yielding a significantly better replication of the initial opening cone in comparison to χ_p .




Figure 4.13: Time evolution of total in-plane spin magnitude S_\perp in the coprecessing frame for the same fiducial binary as in Figure 4.2 (orange). The purple graph shows the spin evolution obtained after applying the χ_p -mapping at the initial time t_0 , the teal graph that of the $\vec{\chi}_\perp$ -parameterisation. We see that $\vec{\chi}_\perp$ preserves the total in-plane spin magnitude, and thus the spin dynamics, much better than the χ_p -mapping.

$\vec{\chi}_\perp$ -mapping represents the full-spin dynamics much more faithfully than χ_p .

As a third diagnostic, we investigate how faithfully both mappings reproduce the

fully spinning precession dynamics, which is represented by the unit quaternions \hat{q}_i . In Figure 4.14 we show the time evolution of the four unit quaternion components of the fiducial binary. The new effective spin mapping $\vec{\chi}_\perp$ clearly replicates the time evolution of each quaternion component much more accurately than χ_p , with the most dramatic improvement observed for \hat{q}_1 and \hat{q}_2 .




Figure 4.14: Time evolution of the quaternion components \hat{q}_0 and \hat{q}_3 (left), and \hat{q}_1 and \hat{q}_2 (right) for the fiducial precessing binary. The system parameterised by $\vec{\chi}_\perp$ (teal) reproduces the precession dynamics of the fully spinning system (orange) much more faithfully than the χ_p -mapping (purple).

To quantify the improvement in replicating the precession dynamics, we perform a match calculation for each of the four quaternion components \hat{q}_i , $i \in [0, 1, 2, 3]$, similar to the white match calculation in Eq. (4.12) but replacing the waveforms h and h_σ with the quaternion components,

$$\mathcal{M}_{\hat{q}_i} = \mathcal{M}(\hat{q}_i, \hat{q}_{\sigma,i}), \quad (4.20)$$

where we use $S_n = 1$, \hat{q}_i denotes the quaternion component from the fully precessing system, and $\hat{q}_{\sigma,i}$ is the quaternion component produced by the effective mapped system, with $\sigma \in [\chi_p, \vec{\chi}_\perp]$.

We compute the quaternion matches for the same 20,833 binaries used in the sky and polarization-averaged strain match calculations. Figure 4.15 shows the results for \hat{q}_1 (left) and \hat{q}_2 (right), which show the largest improvements: The percentage of matches below

0.99 improves from 98.65% with χ_p to 46.33% with $\vec{\chi}_\perp$ for \hat{q}_1 , and for \hat{q}_2 it improves from 95.71% to 46.37%. We see a negligible difference in the results for \hat{q}_0 , which is well reproduced by both spin mappings: None of cases have a match value below 0.99. We see a small improvement in the results for \hat{q}_3 , with the percentage of quaternion matches below 0.99 dropping from 40.5% for χ_p to 34.21% for $\vec{\chi}_\perp$. These results indicate that the observed improvements when using $\vec{\chi}_\perp$ can indeed be attributed to a more faithful representation of precession dynamics itself.

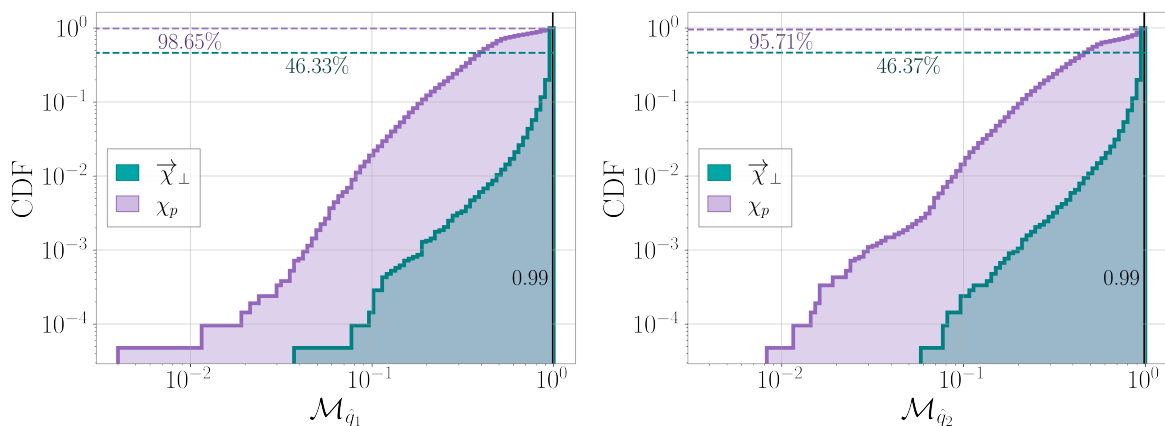


Figure 4.15: Cumulative distribution of matches for two of the four quaternion elements \hat{q}_1 (left) and \hat{q}_2 (right), between the fully precessing dynamics and each of the $\vec{\chi}_\perp$ (teal) and χ_p -mapped (purple) systems. The dashed horizontal lines indicate the percentages of matches which are below 0.99 for each effective mapping. We see significant improvements for the $\vec{\chi}_\perp$ -parameterisation over χ_p .

4.3.3 Accuracy of the Final Spin and Recoil

Finally, we quantify how well the $\vec{\chi}_\perp$ -parameterisation is able to reproduce the final spin and recoil of the remnant black hole. The final mass and spin of the remnant determine the quasinormal modes of the ringdown [177–180], so it is therefore crucial to understand the accuracy with which the final state can be replicated by the reduced set of spin parameters. We will focus on the final spin estimates as previous comparisons against NR simulations have shown that the final mass estimate is only very weakly dependent on precession [127].

To evaluate the final spin using the surrogate model we first evolve the BH spins in the inertial frame from t_0 to a time $t = -100M$ before merger, which are then used to evaluate the remnant fits of Ref. [181] via the public PYTHON package `surfinBH` [182]. The same procedure is followed to obtain the results under the two-spin parameterisations, where either effective spin map is applied at the initial time t_0 . We evaluate the remnant spin for the 20,833 binary configuration of column two in Table 4.2. We assess the accuracy of the final state under the two mappings by calculating the differences in the remnant spin magnitude $\Delta\chi_f$, the final spin tilt angle $\Delta\theta_f$, the azimuthal spin angle $\Delta\phi_f$, the recoil velocity Δv_f and its tilt angle $\Delta\theta_{v_f}$ defined as

$$\Delta\chi_f = \|\vec{\chi}_f\| - \|\vec{\chi}_{f\sigma}\|, \quad (4.21)$$

$$\Delta\theta_f = \|\arccos(\hat{z} \cdot \hat{\chi}_f) - \arccos(\hat{z} \cdot \hat{\chi}_{f\sigma})\|, \quad (4.22)$$

$$\Delta\phi_f = \arccos(\hat{\chi}_{f\perp} \cdot \hat{\chi}_{f\sigma\perp}), \quad (4.23)$$

$$\Delta v_f = \|\vec{v}_f\| - \|\vec{v}_{f\sigma}\|, \quad (4.24)$$

$$\Delta\theta_{v_f} = \|\arccos(\hat{z} \cdot \hat{v}_f) - \arccos(\hat{z} \cdot \hat{v}_{f\sigma})\|, \quad (4.25)$$

where $\sigma \in [\chi_p, \vec{\chi}_\perp]$, and $\vec{\chi}_{f\perp}$ indicates the xy -components of the final spin vector in the inertial frame. We note that the remnant spin and recoil velocities are also returned in the inertial coordinate frame of the NR surrogate, which has no particular physical meaning postmerger. However, as we are computing relative differences in magnitudes and angles, this gauge choice has no effect on the results presented here.

We find marginal improvements in the accuracy of the final spin magnitude and tilt angle using the $\vec{\chi}_\perp$ -mapping as opposed to χ_p . The median tilt angle difference $\Delta\theta_f$ improves slightly from 2.30×10^{-3} rad with χ_p to 1.43×10^{-3} rad with $\vec{\chi}_\perp$; the absolute value $\Delta\chi_f$ also improves slightly from 1.75×10^{-3} for χ_p to 9.96×10^{-4} for $\vec{\chi}_\perp$. However, the largest improvement is found for the azimuthal angle $\Delta\phi_f$, which encapsulates the difference in the relative angle in the xy -plane of the inertial frame as shown in Figure 4.16. We see a dramatic difference between the two mappings, with $\vec{\chi}_\perp$ effectively reproducing

the azimuthal orientation with a median error of less than 0.1 rad, whereas the χ_p -mapping poorly replicates the azimuthal orientation with a median difference of more than 1 rad, and a significant proportion of differences around $\Delta\phi = \pi$. We also note the significantly long tail of the $\vec{\chi}_\perp$ histogram toward angle differences of zero.




Figure 4.16: Error in the azimuthal angle of the final spin state, $\Delta\phi_f$, in radians, between the final spin state produced by the fully precessing waveform, χ_f , and the resulting final spin state of the waveform produced by the $\vec{\chi}_\perp$ (teal) and χ_p (purple) mappings. The new effective spin $\vec{\chi}_\perp$ reproduces the azimuthal angle of the remnant spin much more accurately, reducing the median error to less than 0.1 rad.

We now analyze the effect of the two mappings on the recoil velocity v_f of the final black hole. For the recoil velocity tilt angle $\Delta\theta_{v_f}$, i.e. the polar direction of the recoil, we find a large improvement from a median error of 0.67 rad for χ_p to 0.08 rad for $\vec{\chi}_\perp$ as shown in the right panel of Figure 4.17. For the recoil velocity itself, we only find a modest improvement in Δv_f from a median error of $3.81 \times 10^{-4} c$ for χ_p to $1.47 \times 10^{-4} c$ for $\vec{\chi}_\perp$, corresponding to an improvement in accuracy of ~ 70 km/s on average.

To summarize, overall the $\vec{\chi}_\perp$ -parameterisations reproduce the final state, in particular the orientation of the final spin and the direction of the recoil, much more accurately. Both effective spin parameterisations perform similarly in determining the final spin magnitude.

We note that the comparison using χ_p is not directly comparable to the definition of the final spin used in semi-analytical waveform models, which use (a variety of) in-plane

spin corrections to modify the final spin of an aligned-spin binary [168]. The current generation of precessing phenom models [19, 165, 167] apply a correction of the form S_p/M_f^2 , where M_f is the remnant mass and S_p is an effective in-plane spin contribution. In [165, 167], S_p is taken to be defined as in Eq. (3.14), which is similar to the results presented here obtained by applying the χ_p -mapping [127, 145]. For the more recent model presented in [168], a range of different final spin mappings have been implemented including the χ_p -mapping as well as a precession-averaged mapping that attempts to account for the change in the aligned-spin components due to nutation effects. The EOB models [48] employ the final spin fits of [183] which introduce corrections to the aligned-spin final state fits that depend on the angle between the two in-plane spin vectors and the projection of the spins along the orbital angular momentum. As discussed in [48, 178] and above, a crucial choice is the separation at which the spins are used to evaluate the final state, taken to be $r = 10M$ in [48]. This approach enables EOB models to account for the evolution of the spin vectors ensuring that the same waveform is produced irrespective of the initial separation.




Figure 4.17: Absolute value of the error in the recoil velocity magnitude Δv_f (left) in units of c , and recoil velocity tilt angle $\Delta \theta_{v_f}$ (right) in radians, between the fully spinning waveform and each of the $\vec{\chi}_\perp$ - (teal) χ_p - (purple) mappings. The dashed vertical lines indicate the median error values.

4.4 Discussion

The inclusion of fully relativistic precession effects in semi-analytic IMR waveform models in the strong-field regime remains a challenging problem, with none of the current waveform models from either the phenom or the EOB waveform family including calibration to NR in the precessing sector. The high dimensionality of the precessing BBH parameter space obfuscates a clear path for calibration. Effective spin parameterisations to reduce the number of spin degrees of freedom are a promising way forward to including fully relativistic precession in the strong-field regime. Previously, a scalar quantity χ_p was introduced to this effect but its efficacy was only demonstrated for the inspiral regime [174]. Here, we have assessed its applicability in the strong-field regime. Crucially, we have found that while χ_p does faithfully represent the (2,2)-mode of the majority of fully precessing systems, it does not accurately reproduce HMs. Since HMs are excited by mass and spin asymmetries, which can be very pronounced for precessing binaries, they carry crucial parameter degeneracy breaking power [95, 97, 98, 102, 108, 115] making the accurate modelling of HMs critical. Therefore, NR calibration through a simple χ_p -parameterisation is unlikely to be sufficient to satisfy the accuracy requirements for future GW observations.

To improve upon the shortcomings of χ_p , we have introduced a new two-dimensional effective precession spin vector, $\vec{\chi}_\perp$, and have performed extensive studies comparing the efficacy of $\vec{\chi}_\perp$ to that of χ_p in the strong-field regime using the NR surrogate waveform model NRSur7dq4 [160]. When analyzing individual $h_{\ell m}$ -modes, in particular the (2, ± 1)-modes, we have found that $\vec{\chi}_\perp$ performs significantly better than χ_p , but both effective parameterisations yield comparable results for the quadrupolar (2, ± 2)-modes. Correspondingly, we also have found a significant improvement in the precessing strain matches with the new mapping, from which we have concluded that the improved efficacy of $\vec{\chi}_\perp$ over χ_p for the HMs has a significant effect on the accuracy of the overall strain, demonstrating the importance of accurately modelling HMs in precessing systems. Furthermore, we have found that $\vec{\chi}_\perp$ performs better compared to χ_p in the equal-mass limit

(see Figure 4.6). In this limit, the BH spins precess at the same rate, locked in orientation relative to each other. The parameter χ_p , which is defined to mimic the average rate of precession, performs knowingly poorly in this limit [174]; $\vec{\chi}_\perp$, on the other hand, is constructed such that it approximates the total in-plane spin of the fully precessing system at some reference time, leading to a significantly improved behavior in the equal-mass limit as anticipated from PN theory [14, 184]. As expected, we have found that $\vec{\chi}_\perp$ performs increasingly better for larger mass ratios $q \gtrsim 2$, where the spin on the smaller BH becomes negligible and hence the approximation with a single in-plane spin becomes more accurate. However, in the intermediate region between these two regimes, whilst still a considerable improvement upon χ_p , we have found a small drop in accuracy at a mass ratio of $q \sim 1.4$, where two-spin effects are important but are not fully captured in $\vec{\chi}_\perp$.

We have further demonstrated that the overall improvement relative to a χ_p -parameterisation can be attributed to a more accurate replication of the precession dynamics itself when using the $\vec{\chi}_\perp$ -parameterisation. Indeed, in the case where only one of the two objects has nonzero in-plane spin components, the full dynamics are returned exactly, which is not the case for χ_p . The accurate capture of the precession dynamics of particular interest as a natural way for incorporating strong-field precession information into waveform models is through calibrating the precession dynamics itself, i.e. the rotation operator \mathbf{R} of Eq. (3.12) or, equivalently, the quaternions.

Additionally, we have also quantified how well the $\vec{\chi}_\perp$ -mapping is able to replicate the final spin and recoil velocity of the remnant black hole. We have found a considerable improvement in the accuracy with which we have replicated the azimuthal direction of the remnant spin, and moderate improvements in the accuracy of the magnitude and direction of the recoil velocity. This suggests our $\vec{\chi}_\perp$ -mapping is better able to replicate the final direction of GW emission, compared to χ_p . Previous work has demonstrated that the relative orientation of the in-plane spins at merger plays a crucial role in determining the final state properties [185–188]. We have attributed the observed improvements to

the (partial) incorporation of two-spin effects, which are crucial for determining the recoil direction and velocity of the final BH.

Despite its significantly better performance in all areas, there are also caveats associated with $\vec{\chi}_\perp$: (i) For spin configurations with similar in-plane spin magnitudes, i.e. $S_{1\perp} \simeq S_{2\perp}$, we expect larger mismatches due to, by construction, the neglect of larger in-plane spin-spin couplings. (ii) We have normalized $\vec{\chi}_\perp$ such that the Kerr limit is not violated. Consequently, for binaries with large spin magnitudes, $\vec{\chi}_\perp$ will underestimate the magnitude of the in-plane spin in the system. Due to the limited spin parameter range of the surrogate, we have not been able to fully quantify the effect of this on the performance of the mapping. (iii) The conditional placement of $\vec{\chi}_\perp$ on either of the two black holes introduces a discontinuity, in the sense that waveforms with $\vec{\chi}_\perp$ placed on the primary BH show slightly different features from those with $\vec{\chi}_\perp$ on the secondary BH. We note that all of our $\vec{\chi}_\perp$ -mapped individual waveforms are physical and continuous, but that a shift in phenomenological features can occur between binary configurations with $S_{1\perp} = S_{2\perp} + \epsilon$ where $\vec{\chi}_\perp$ is placed on the primary BH, against the same binary configuration with slightly smaller $S_{1\perp} = S_{2\perp} - \epsilon$, where $\vec{\chi}_\perp$ will be placed on the secondary BH. With this in mind, we have tested the performance of $\vec{\chi}_\perp$ without the conditional placement. We have recalculated the sky-and-polarization-averaged strain matches shown in Figure 4.10 with $\vec{\chi}_\perp$ always placed upon the primary BH irrespective of whether the precession is dominated by the primary or secondary BH, and indeed have found little difference from the original $\vec{\chi}_\perp$ strain match distribution, with the median mismatch increasing minimally from 2×10^{-4} to 2.08×10^{-4} . Additionally, we have recalculated the white noise mode-by-mode matches of Eq. (4.12), again with $\vec{\chi}_\perp$ always placed on the primary BH. While we have found little difference between the $(2, \pm 2)$ -mode results for $\vec{\chi}_\perp$ with and without conditional placement, we have found that it has a marked effect on the results for HMs. For example, in the $(2, 1)$ -mode at mass ratio $q = 3$, the percentage of mismatches below 0.99 using $\vec{\chi}_\perp$ without conditional placement rises to 41.9%, compared to just 9.1% if we include the conditional placement (under the χ_p -mapping the value is

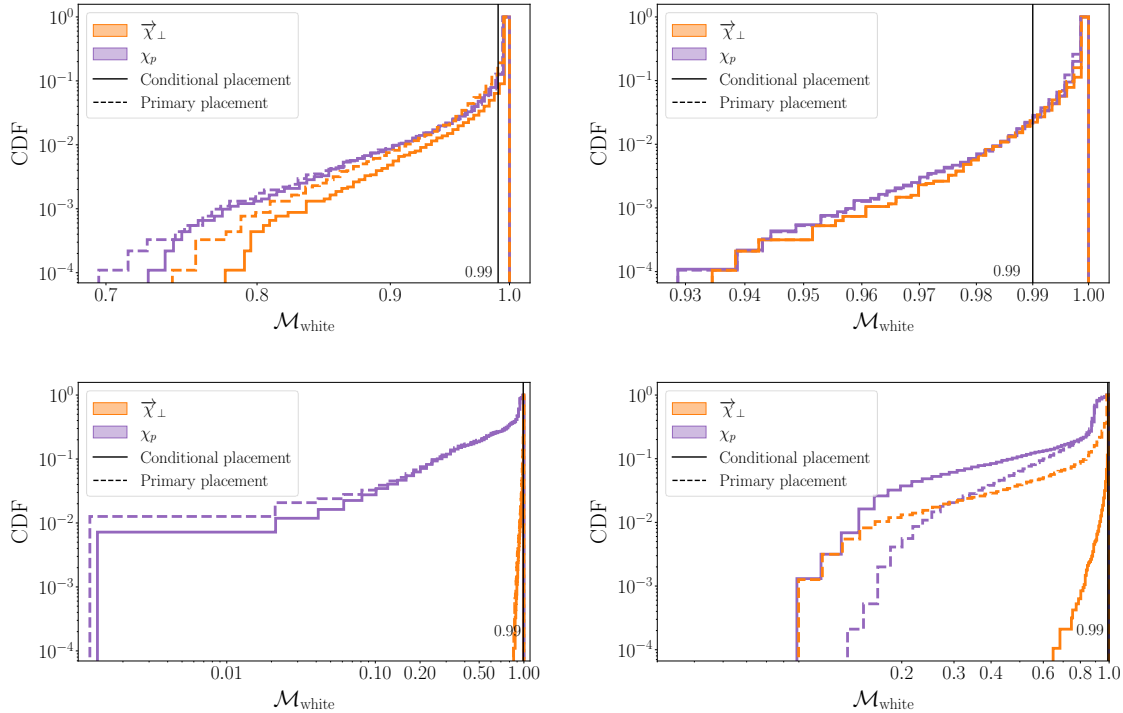


Figure 4.18: Cumulative histograms of white noise mode-by-mode matches for the (2,2)-mode (top row) and the (2,1)-mode (bottom row) for mass ratios $q = 1$ (left column) and $q = 3$ (right column) for the same binaries as in Figure 4.6 with the $\vec{\chi}_{\perp}$ -parameterisation (orange) and the χ_p -parameterisation (purple). The solid outlines represent the parameterisations including conditional placement, whereas the dashed lines show results when the effective spin is always placed on the primary black hole. The effect of conditional placement is most noticeable at $q = 3$ in the (2,1)-mode, where $\vec{\chi}_{\perp}$ with conditional placement dramatically outperforms other mappings. We also note that including conditional placement improves the performance of both effective spin parameterisations in most cases, although this improvement is negligible for the (2,2)-mode at mass ratio $q = 3$. The exception to this is the χ_p -parameterisation for the (2,1)-mode at mass ratio $q = 3$, where conditional placement worsens the match distribution for χ_p .

98.8%). We therefore have concluded that for HMs, it is crucial to accurately capture spin asymmetries by placing the effective spin appropriately, to achieve an accurate mapped waveform mode.

Lastly, we have also tested whether the improvements found by using $\vec{\chi}_{\perp}$ over χ_p are entirely due to the conditional placement, and whether an analogous conditional placement of χ_p would have similar effects. In addition to the individual case shown

in Fig. 4.3, in Fig. 4.18 we recalculate the white noise mode-by-mode matches shown in Fig. 4.6, for both spin parameterisations χ_p and $\vec{\chi}_\perp$, using (i) conditional placement (solid) and (ii) placement always on the primary black hole (dashed). We note that in all four panels, the best performance is obtained when conditionally placing $\vec{\chi}_\perp$. The (2,2)-mode at mass ratio $q = 1$ (top left) shows a small improvement in both parameterisations' performance when conditional placement is included. The biggest improvement can be seen in the (2,1)-mode at mass ratio $q = 3$ (bottom right), where a conditionally placed $\vec{\chi}_\perp$ dramatically outperforms all other configurations, and neither a conditionally placed χ_p , nor a $\vec{\chi}_\perp$ affixed to the primary would be able to achieve these improvements. Interestingly, this panel also displays the only instance where conditional placement can worsen the performance of χ_p . Therefore, we conclude that particularly for HMs at unequal-mass ratios, to obtain the dramatic improvements we have seen, both the new effective spin $\vec{\chi}_\perp$ and conditional placement are required.

An example of imposing this condition also on χ_p is shown in Figure 4.3 for the fiducial binary, and we have indeed seen that the phenomenology is captured better. To quantify the improvement in the performance of χ_p when imposing conditional placement, we have recalculated the sky-and-polarization-averaged strain matches shown in Figure 4.10 with an analogous conditional placement for χ_p . We have found only a small improvement compared to the χ_p -mapping without conditional placement, with the median strain mismatch improving from 4×10^{-3} to 3.4×10^{-3} , compared to a median of 2×10^{-4} with $\vec{\chi}_\perp$. We have also recalculated the mode-by-mode white noise matches shown in Figure 4.6, for both effective spin parameterisations, with and without conditional placement. These results are shown in Figure 4.18. We see that for HMs at unequal-mass ratios, neither a conditionally placed χ_p , nor $\vec{\chi}_\perp$ always placed on the primary object, can replicate the dramatic improvements we have previously seen in Figure 4.6. We therefore have surmised that the improvements we have seen in the performance of $\vec{\chi}_\perp$ over χ_p are due to a combination of both the new parameterisation itself, and the conditional placement, and that both are a necessary requirement to reproduce accurate precessing higher-order

waveform modes.

Finally, we also note that the efficacy of $\vec{\chi}_\perp$ has not been investigated for the special case of transitional precession, which leads to the tumbling of the total angular momentum \hat{J} when $L \simeq S$ and $\hat{L} = -\hat{S}$. As with all effective mapping that neglects some spin contributions, however, we expect that the fine-tuned conditions needed for the occurrence of the transitional precession phase are not preserved under the mapping.

In conclusion, our results have demonstrated that by introducing the two-dimensional vector quantity $\vec{\chi}_\perp$, which partially accounts for two-spin effects, we can accurately reproduce the waveforms of fully precessing binaries, and in particular their HMs, in the strong-field regime across a wide range of the BBH parameter space. The effective reduction of four in-plane spin components to two provides a clear and tractable path forward to meaningfully incorporating precession effects in the strong-field regime into semi-analytic waveform models with HMs, which we leave to future work.

5. Accelerating Multimodal Gravitational Waveform Models from Precessing Compact Binaries with Artificial Neural Networks

This chapter is a reformatted version of work presented in Ref. [40], "Accelerating Multimodal Gravitational Waveforms from Precessing Compact Binaries with Artificial Neural Networks", published in *Phys. Rev. D* 106.104029, and of which L. M. Thomas is the lead author. The idea to produce a neural network surrogate model for a precessing, higher mode waveform model was proposed by G. Pratten and P. Schmidt, building upon work in Ref. [153] where a surrogate model for the quadrupolar mode of aligned-spin BBH waveform model `SEOBNRv4` was constructed using neural networks. All of the training, validation and test datasets as described in Subsec. 5.2.1 were produced by L. M. Thomas, including both the coprecessing frame mode and Euler angle data. She constructed the coprecessing frame modes portion of the model, including the reduced bases and empirical interpolants as described in Subsec. 5.2.2, and parameter space fitting with neural networks as described in Subsec. 5.2.3. L. M. Thomas also produced the coprecessing frame mode mismatches as shown in Figs. 5.6. G. Pratten constructed the Euler angle portion of the model, including reduced bases, empirical interpolants and parameter space fitting. He also produced the code to twist up the predicted coprecessing waveforms using the predicted Euler angles to produce the fully precessing strain, and calculated the full precessing strain mismatches. G. Pratten also carried out the model timing tests in Subsec. 5.3.2. The discussion and interpretation of our model results was a joint effort between L. M. Thomas, G. Pratten and P. Schmidt. All the Figs. in this chapter were produced by L. M. Thomas except for Figs. 5.5, 5.10, 5.11, 5.12 and 5.13. All three authors contributed to the original text with thesis-specific expansions and additions made by L. M. Thomas.

In Chapter 4, we discussed the problem of incorporating strong-field precession effects

into semi-analytic waveform models through precessing NR calibration. We introduced a new effective precession spin $\vec{\chi}_\perp$, which replicates precessing morphology, notably in the higher modes, with fewer spin degrees of freedom. In this chapter, we tackle an equally important but contrasting aspect of waveform modelling, namely that time-domain semi-analytic models, and in particular those which include more physical effects for completeness such as precession effects, can be very slow to evaluate. This means that in practice, they can sometimes be unusable in real parameter estimation scenarios without the use of reduced order modelling techniques. In this work we combine reduced-order modelling with the power of artificial neural networks (ANNs) to build a computationally vastly more efficient surrogate model of the state-of-the-art inspiral-merger-ringdown (IMR) waveform model `SEOBNRv4PHM` [48] that includes both spin-induced orbital precession [14] and higher-order modes beyond the quadrupole emission. While the efficacy of this approach has previously been demonstrated for the quadrupole $((2, 2)-)$ mode of aligned-spin binary black holes (BBHs) [153, 154], here we demonstrate its feasibility for the multimodal, precessing case. To achieve this, we decompose the `SEOBNRv4PHM` waveform model into eight components that describe the modes in a non-inertial, co-precessing coordinate frame and three components that encode the precession dynamics. Using a combination of traditional surrogate modelling steps and neural networks to produce parameter fits, we build a fast surrogate model for each component. Using extensive optimisation we determine an optimal network for each component, which allows us to speed up the model evaluation by a factor of a few hundred on average on a CPU and even further on a graphics processing units (GPU), demonstrating the efficacy of this approach for state-of-art multimodal waveforms with precession.

The chapter is organised as follows: First in Sec. 5.1 we introduce the methodology behind this work, including a brief overview of artificial neural networks in Sec. 5.1.1, and details of precessing waveform decomposition in Sec. 5.1.2. We then detail the construction of the model in Sec. 5.2, first describing the training data upon which the model is built in Sec. 5.2.1, the reduced basis and empirical interpolant construction in Sec. 5.2.2, and the

neural networks for each of the model components in Sec. 5.2.3, both coprocessing modes and Euler angles, putting these together to describe the full surrogate model construction in Sec. 5.2.4. We then discuss the completed model in Sec. 5.3, both the accuracy of the waveforms generated in Sec. 5.3.1 and the timing of the model evaluation in Sec. 5.3.2. Finally, we summarise the model construction and results in Sec. 5.4, also discussing caveats and further work.

5.1 Methodology

In this section, we provide details of the key ingredients required to build our ANN surrogate model. As previously mentioned in Sec. 3.5, surrogate models are fast, accurate approximations to an underlying (slower) model, over a chosen parameter space region. After building a reduced basis and empirical interpolant, a parameter space fit needs to be constructed to fit the reduced basis coefficients at the EI nodes to the intrinsic binary parameters. In this work, we will follow Ref. [153] and use ANNs to obtain the fitting coefficients. To construct the reduced bases and empirical interpolants we use the publicly available Python package RomPy [152, 189].

In Subsec. 5.1.1 we now outline ANNs as relevant to this thesis, including their training and evaluation for use in a surrogate model, and then in Subsec. 5.1.2 we describe the decomposition of our precessing waveform into components which we then model separately.

5.1.1 Artificial Neural Networks

An artificial neural network (ANN) [190] is a computational system inspired by biological neural networks which make up the brain. They have become widespread tools for a wide variety of uses such as classification, clustering, function approximation or regression, to name a few. Here we give only a brief overview as is relevant to this work, but for a thorough overview of the subject we refer the reader to, eg. Ref. [191]. An ANN comprises

a set of connected neurons (or nodes) between which information is to be passed, along edges which connect them. The nodes are typically arranged in layers, such that the input (a set of real numbers) traverses and is transformed through subsequent layers of nodes before being output. This process of passing data through the ANN from input to output is called forward propagation. If an ANN is fully connected, then each node in a particular layer will have an edge which connects it to each node in the previous layer. Each node will have an associated nonlinear activation function, which transforms the input, and each node and edge will have weights which adjust as the network is trained. To be more concrete, let us consider a particular neuron within the network with an index j . Assuming the ANN is fully connected, this node will receive one real number input from each node in the previous layer, x_i . These inputs are multiplied by a set of corresponding weights w_{ij} , where there is a different set of weights for each node, and a bias b_j may be added such that the inputs become

$$x_i \Rightarrow x_i w_{ij} + b_j. \quad (5.1)$$

The transformation thus far has been entirely linear, but now this value is passed to the nonlinear activation function $f(x)$ of the node. An example is the **Softplus** activation function which is given by

$$f(x) = \log(1 + \exp(x)). \quad (5.2)$$

Therefore the output of this node, x_j is given by

$$x_j = f(x_i w_{ij} + b_j). \quad (5.3)$$

For best performance, an ANN must be trained, or equivalently, the weights and biases must be fine-tuned. This process requires a set of training data, made up of inputs and outputs to the neural network, where the inputs are passed through the network many times, each iteration called a training epoch. With each epoch the outputs predicted

by the neural network and the expected outputs from the training dataset are compared using a metric (often the mean squared error) and the weights and biases updated using an optimisation algorithm (optimiser). This is known as supervised training, because with each epoch the predicted and expected outputs are compared to guide the neural network learning process. The learning rate is a hyperparameter of the neural network which controls how much the weights and biases are changed by the optimiser at the end of the epoch. Therefore prior to training, the user chooses the neural network hyperparameters such as the architecture, which include the number and size of layers, the activation function, optimiser, and learning rate. A training set must then be specified and used to train the ANN for a set number of epochs, or until the network reaches an accuracy condition upon the output predictions.

5.1.2 Waveform Decomposition

As previously described in Chap. 2, binary black holes on quasi-spherical orbits span a seven-dimensional (intrinsic) parameter space characterised by the mass ratio $q = m_1/m_2 \geq 1$ and the (dimensionless) spin angular momenta $\vec{\chi}_1$ and $\vec{\chi}_2$. If the spin angular momenta are misaligned with the direction of the instantaneous orbital angular momentum $\hat{L}(t)$, then spin-induced precession occurs [14, 15]. This causes the orbital plane to change its spatial orientation as the binary inspirals due to GW emission. This more complex two-body dynamics leads to amplitude and phase modulations of the emitted GW signal $h(t; \vec{\lambda})$ and is also a source of the excitation of higher-order multipoles, $h_{\ell m}$ (see Eq. (3.3)), in the radiation field, which must be included to accurately describe the GW signal. Due to the increased complexity, modelling the signal from precessing BBHs is a challenging task but is accomplished as follows [14, 123, 128, 173]: The GW modes from precessing binaries, $h_{\ell m}^P(t; \vec{\lambda})$, can be conveniently decomposed into a simpler carrier signal corresponding to a non-inertial coprecessing observer, $h_{\ell m}^{\text{coprec}}(t; \vec{\lambda})$, and a time-dependent rotation operator \mathbf{R} which encodes the orbital precession dynamics, as seen in Eq. (3.12).

As a first approximation, the coprecessing waveform modes can be approximated

by aligned-spin modes [128, 173]. This simplifying approximation is made in many of the state-of-the-art waveform models [48, 165, 168] and is a known source of modelling errors [120]. Importantly, this approximation assumes a conjugate symmetry between the $+m$ and $-m$ modes, which no longer holds in the case of precessing binaries [128, 131]. The waveform model we emulate here, SEONBNRv4PHM [48, 111], contains the $(2, \pm 2), (2, \pm 1), (3, \pm 3), (4, \pm 4)$ and $(5, \pm 5)$ coprecessing modes defined in a time-dependent coordinate frame that tracks $\hat{L}(t)$ (L -frame), and assumes conjugate mode symmetry, i.e.,

$$h_{\ell, -m}^{\text{coprec}}(t; \vec{\lambda}) = (-1)^\ell h_{\ell m}^{\text{coprec}*}(t; \vec{\lambda}). \quad (5.4)$$

Therefore, we only model the positive m -modes in the coprecessing frame and obtain the $-m$ -modes via Eq. (5.4). The coprecessing waveform modes are then further decomposed into amplitude and phase,

$$h_{\ell m}^{\text{coprec}}(t; \vec{\lambda}) = A_{\ell m}(t; \vec{\lambda}) e^{i\phi_{\ell m}(t; \vec{\lambda})}. \quad (5.5)$$

For the rotation operator we will use its $\text{SO}(3)$ representation and model the three Euler angles $\alpha(t; \vec{\lambda}), \beta(t; \vec{\lambda})$ and $\gamma(t; \vec{\lambda})$ in an inertial Cartesian coordinate frame that is aligned with the total angular momentum at the initial time t_0 , i.e. $\mathbf{J}(t_0) = \hat{z}$, as shown in Fig. 3.3, henceforth referred to as the J -frame.

Utilising this decomposition, we build (i) ANN surrogates for the amplitude and phase of each coprecessing positive m -mode contained in SEONBNRv4PHM (see Sec. 5.2.3) and (ii) ANN surrogates for the three Euler angles (see Sec. 5.2.3).

5.2 Model

In this section we describe the construction of our surrogate model. As described in Sec. 5.1.2, we model the coprecessing modes and Euler angles separately. We detail the training dataset upon which the model is built in Sec. 5.2.1, as well as the validation

Training data subset		Number of binaries	q	$ \chi_1 $	θ_1 [rad]	ϕ_1 [rad]
Systematically sampled	Non-spinning	6	[1, 1.2, 1.4, 1.6, 1.8, 2]	[0]	-	-
	Spin-aligned	48	[1, 1.2, 1.4, 1.6, 1.8, 2]	[0.2, 0.4, 0.6, 0.8]	$[0, \pi]$	-
	Precessing	720	[1, 1.2, 1.4, 1.6, 1.8, 2]	[0.2, 0.4, 0.6, 0.8]	$[\pi/6, \pi/3, \pi/2, 2\pi/3, 5\pi/6]$	$[0, \pi/3, 2\pi/3, \pi, 4\pi/3, 5\pi/3]$
Randomly sampled		199, 226	$U [1, 2]$	$U [0, 0.8]$	$U [0, \pi]$	$U [0, 2\pi]$

Table 5.1: Parameters of the 200,000 binaries which span our training dataset. The training space is split into two sectors: (i) a systematically sampled subset which is included to ensure coverage of the parameter space boundaries; (ii) a random but uniformly sampled subset. All spin parameters are quoted at a reference frequency of 4Hz for the (2, 2)-mode for a binary with a total mass of $60M_\odot$

and test datasets. In Sec. 5.2.2 we describe the construction of the reduced bases and empirical interpolants for each component, and in Secs. 5.2.3 and 5.2.3 we describe the neural network architecture and training for the coprecessing modes and Euler angles respectively, assessing the accuracy of each component. We will then consider the accuracy and speed of the complete model to produce a fully precessing signal in Sec. 5.3.

5.2.1 Training, Validation and Testing Data

Waveforms

Our waveform training dataset consists of 2×10^5 multipolar SEOBNRv4PHM waveforms with mass ratios $q \in [1, 2]$ and arbitrarily oriented spin on the primary black hole with magnitude $|\chi_1| \leq 0.8^1$; the secondary black hole is nonspinning. Both the coprecessing waveform modes and the time-dependent Euler angles are obtained directly from the

¹We note that the spin orientation is defined relative to the orbital angular momentum $\hat{L}(t_0)$ at the initial time. Decomposed in Cartesian coordinates the spin vector is given by $\{\chi_{1x}, \chi_{1y}, \chi_{1z}\}$, where $\chi_{1z} = \vec{\chi}_1 \cdot \hat{L}(t_0)$.

SEOBNRv4PHM implementation in the public LIGO Algorithm Library LAL [49], mitigating the need to perform any additional post-processing.

We first randomly sample 199,226 binaries from the reduced parameter space, drawing the parameters from distributions uniform in mass ratio q , uniform in spin magnitude $|\chi_1|$ and isotropic in spin orientation (θ_1, ϕ_1) . We supplement these binaries with an additional 774 systematically chosen points to accurately sample the boundaries. The parameters of these systematically chosen binaries are listed in Tab. 5.1 and a visualisation of the training set can be found in Fig. 5.1, which shows the distribution of the 2×10^5 waveform training dataset, plotted in the space of primary spin components and coloured by number density.




Figure 5.1: Visualisation of the spin parameters of the entire training dataset coloured by number density.

We initially generate waveforms such that the $(2, 2)$ -mode of a binary with a total mass of $60 M_\odot$ starts from an initial frequency of 4Hz. In geometrised units this corresponds to an approximate length of $\sim 2 \times 10^5 M$ before merger, though the duration varies due to mass ratio and inspiral spin [76, 192]. The modes are then aligned such that the peak of the quadrature of all modes occurs at $t = 0M$. Each waveform is first generated on a uniform time grid with a time spacing $\Delta t = 0.1M$, and then reinterpolated onto a non-uniform grid which is 20 times coarser in the early inspiral, but retains the $0.1M$ spacing in the later-inspiral, merger and ringdown. The waveforms are of varying length pre- and

post-merger, and so have different time grids of the same resolution, but are required to be of equal length and evaluated upon the same times in order to build the reduced bases. As a data preprocessing step we choose the waveform with the shortest length pre-merger, and reinterpolate all waveforms onto this common time grid. We then choose the waveform with the shortest post-merger signal and truncate all waveforms such that the final time matches that of the shortest. We note that this truncation is less than $10M$ for all waveforms, and contains a negligible amount of the ringdown signal in all cases. For computational reasons, we truncate all waveforms to be of length 10^4M pre-merger. Therefore all waveforms start at 10^4M before the peak, and include $110M$ of post-merger signal. We note that due to this truncation, the spin parameters are specified at the initial time $t \sim -2 \times 10^5M$ and not at the start of the waveforms. Since the spins in precessing binaries evolve with time, it is necessary to define the reference time or frequency at which they are defined. Being able to do this for some arbitrary time/frequency requires either code (see e.g. [114]) or additional NNs that track the spin evolution. We leave building neural networks for the spin evolution for future work. The truncated waveforms are then re-interpolated onto a uniform grid with spacing $\Delta t = 1M$ in order to build the reduced bases and empirical interpolants, as we found that the finer $0.1M$ spacing was not required.

We note that when constructing the models for the coprecessing odd- m mode amplitudes and phases, not all of the 2×10^5 training waveforms are used. We first remove training points where there is very little spin or mass asymmetry in the system, as we expect the odd- m amplitudes to be small and therefore noisy in the true `SEOBNRv4PHM` data. We impose a cut of $q > 1.01$, $\chi_{1z} > 10^{-2}$, which removes 109 points from the training set. Next, we remove any training data which show signs of (unphysical) discontinuities in the phase, possibly due to next-to-quasi-circular corrections in the `SEOBNRv4PHM` data. For the coprecessing (2, 1)-mode, this amounts to 11,091 points, and 198 for the (3, 3)-mode. Therefore, for the (2, 1)-mode amplitude and phase, the total training dataset is 188,800 waveforms, whereas for the (3, 3)-mode it is 199,693.

To validate our neural networks as they train, we also produce a validation dataset of 10^4 waveforms which covers the same parameter space as the training set. We sample this validation set uniformly in mass ratio, primary spin magnitude, spin tilt and azimuthal angles. All preprocessing steps for the validation data are the same as for the training data: we interpolate these waveforms onto the same common time grid with an equivalent spacing. For the coprocessing odd- m modes, we remove 4 points with little asymmetry, 553 which show signs of phase discontinuity in the (2, 1)-mode phase, and 9 in the (3, 3)-mode. This equates to a validation set size of 9, 443 for the coprocessing (2, 1)-mode, and 9, 987 for the (3, 3)-mode.

Lastly, we also produce a separate test dataset of 10^4 waveforms in exactly the same way as the validation set, which is completely independent and unseen by the neural networks. Of this dataset, 3 points are removed for the coprocessing odd- m modes due to little symmetry, 672 due to discontinuities in the (2, 1)-phase, and 13 due to the (3, 3)-phase. Therefore for the (2, 1)-mode amplitude and phase, the test set is of size 9325, and for the (3, 3)-mode it is 9, 984.

Euler Angles

For the Euler angles, we use the same dataset of 2×10^5 waveforms as described above. However, as the Euler angles become ill defined in the non-precessing limit, we restrict our training data to only those binaries with an initial in-plane spin magnitude $|\chi_{1,\perp}| = \sqrt{\chi_{1x}^2 + \chi_{1y}^2} > 10^{-3}$. In contrast to above, we decompose this initial dataset into a training dataset of 1.8×10^5 binaries and a validation dataset of 18, 634 binaries. As no hyperparameter optimization was performed on the Euler angle networks, the validation dataset is never used to train the network or to inform the network hyperparameters. We therefore treat the validation dataset as being effectively independent. The data conditioning is identical to the procedure described above for the waveform modes, with the Euler angles being evaluated on a uniform grid with spacing $\Delta t = 1M$ and a length of 10^4M .

5.2.2 Reduced Basis and Empirical Interpolant

We construct our reduced bases and empirical interpolants following the algorithm described in Sec.3.5. We separate each coprocessing mode into its constituent amplitude and phase, and construct a reduced basis, empirical interpolant and neural network for each component. We also construct a reduced basis and empirical interpolant for each Euler angle separately, leading to a total of 11 different components to make up the full precessing signal². When discussing the construction and evaluation of these models, we use the following terminology: \mathbf{X} describes the input parameters of the model, i.e. the four intrinsic parameters of the binary $\mathbf{X} = \vec{\lambda} = \{q, \chi_{1x}, \chi_{1y}, \chi_{1z}\}$; \mathbf{Y} is an n -dimensional vector that denotes the fitting coefficients, for example the mode amplitudes in Eq. (3.29).

We choose to condition the data before building our reduced bases as we found this to be beneficial for the neural network performance: For the coprocessing modes we use a `scikit-learn` [193] `Standard` scaler on the \mathbf{X} data and a `MinMax` scaler on the \mathbf{Y} data for the phases as we found that without scaling the greedy algorithm for the coprocessing (2,1)- and (3,3)-mode phases was unable to converge and produce a reduced basis to within the greedy tolerance accuracy. We also remove the initial phase at time $t = -10,000M$, such that all phase data begin at zero. We note that we do not explicitly model these initial phases, and leave this to future work.

In contrast, we find no major benefit to scaling the \mathbf{X} data for the Euler angles and the amplitude \mathbf{Y} data for the coprocessing modes. For α and γ , we apply a `MinMax` scaler to the \mathbf{Y} data but we do not apply any preprocessing to the \mathbf{Y} data for β . A summary of the data conditioning can be found in Tab. 5.2. Tab. 5.3 shows the median and maximum mismatches for the reduced basis and empirical interpolant representations for each of the 11 components of the model, across both the training and validation datasets.

To build the reduced bases, we use an absolute greedy error tolerance of $\sigma = 10^{-6}$ for all components of the coprocessing modes, except for the phases of the (2,1) and

²The odd- m modes are obtained via conjugation and hence do not need to be modelled separately but are included in the full precessing signal.

(ℓ, m)	Component	Training Set Size	Greedy Tolerance	Basis Size
(2, 2)	Amplitude	200,000	10^{-6}	23
(2, 2)	Phase		10^{-6}	29
(2, 1)	Amplitude	188,800	10^{-6}	26
(2, 1)	Phase		10^{-8}	40
(3, 3)	Amplitude	199,693	10^{-6}	26
(4, 4)	Phase		10^{-3}	46
(4, 4)	Amplitude	200,000	10^{-6}	4
(4, 4)	Phase		10^{-6}	29
Euler Angle				
	α	180,000	7×10^{-9}	18
	β	180,000	6×10^{-7}	19
	γ	180,000	7×10^{-9}	18

Table 5.2: Greedy tolerances and reduced basis sizes for the amplitude and phase of each mode, as well as the Euler angles. For phases, **MinMax** scaling was used on the Y data. For both amplitudes and phases, standard scaling was used on the X data. For the Euler angles, **MinMax** scaling was used on the Y-data for α and γ .

(ℓ, m)	Component	$\bar{\mathcal{M}}_f^{\max, \text{train}}$	$\bar{\mathcal{M}}_f^{\text{median, train}}$	$\bar{\mathcal{M}}_f^{\max, \text{val}}$	$\bar{\mathcal{M}}_f^{\text{median, val}}$
(2, 2)	Amplitude	1.1×10^{-5}	5.0×10^{-7}	1.4×10^{-5}	6.0×10^{-7}
(2, 2)	Phase				
(2, 1)	Amplitude	3.2×10^{-3}	1.8×10^{-6}	3.4×10^{-4}	1.7×10^{-6}
(2, 1)	Phase				
(3, 3)	Amplitude	1.8×10^{-2}	3.0×10^{-4}	2.7×10^{-2}	3.1×10^{-4}
(3, 3)	Phase				
(4, 4)	Amplitude	1.6×10^{-4}	4.0×10^{-5}	1.9×10^{-4}	4.1×10^{-5}
(4, 4)	Phase				
Euler Angle		$\bar{\mathcal{M}}_t^{\max, \text{train}}$	$\bar{\mathcal{M}}_t^{\text{median, train}}$	$\bar{\mathcal{M}}_t^{\max, \text{val}}$	$\bar{\mathcal{M}}_t^{\text{median, val}}$
	α	3.0×10^{-5}	1.4×10^{-9}	2.6×10^{-6}	1.4×10^{-9}
	β	4.1×10^{-6}	1.3×10^{-6}	4.4×10^{-6}	1.3×10^{-7}
	γ	4.2×10^{-5}	1.5×10^{-9}	1.0×10^{-6}	1.5×10^{-9}

Table 5.3: Maximum and median training and validation dataset mismatches for the amplitude and phase of each mode, as well as the Euler angles. For phases, **MinMax** scaling was used on the Y data. For both amplitudes and phases, standard scaling was used on the X data. For the Euler angles, **MinMax** scaling was used on the Y-data for α and γ and we use the time-domain mismatch $\bar{\mathcal{M}}_t$ as our metric.

(3, 3)-modes. For the (2, 1)-mode, we decreased the greedy tolerance to 10^{-8} as we found a significant tail of poor mismatches against the reduced basis representation with a

tolerance of 10^{-6} . Conversely, for the (3,3)-mode, we reduced the tolerance to 10^{-3} in order to achieve a reduced basis of manageable size. The tolerances and the sizes of the resulting reduced bases (and therefore the number of empirical interpolation nodes) are given in the fourth and fifth column of Table II.

To assess the accuracy of the coprocessing (ℓ, m) -modes reconstructed from their reduced basis representations in amplitude and phase, we compute frequency-domain white noise mode-by-mode mismatches $\bar{\mathcal{M}}_f$, defined by Eq. (2.74) against the original SEOBNRv4PHM data. Columns 6-9 of Table II show the maximum and median mismatch across the full training and validation datasets for each coprocessing mode, noting that the validation data is not used in the construction of the reduced bases. Generally, we find that the odd- m modes are less accurately represented than the even- m modes and that their bases sizes are larger. This is perhaps not too surprising as the odd- m modes are (i) subdominant and (ii) contain more structure, therefore requiring more basis elements to achieve the same representation accuracy [157].

Similarly, we compute time-domain mismatches $\bar{\mathcal{M}}_t$, defined by Eq. (2.78) between the original SEOBNRv4PHM data Euler angles, and those reconstructed from the reduced basis projections. We do this across both the training and validation datasets, and state the median and maximum values for each dataset in columns 5-8 in the bottom half of Table II. We see that for both the coprocessing modes and the Euler angles, the median mismatch across both datasets is comparable to the greedy tolerance used to create the reduced basis (for the mode mismatches, it is limited by whichever greedy tolerance is larger, amplitude or phase).

Lastly, the similar mismatches for the coprocessing modes across both the training dataset, which was used to construct the bases, and the validation dataset, which was previously unseen, suggests that the reduced bases are large enough to accurately represent waveforms across our chosen parameter space. We note that for the Euler angles, the mismatches (both median and maximum) over the validation dataset can be up to an order of magnitude smaller than over the training dataset. This suggests that the validation

dataset is not large enough to accurately represent the full distribution over the entire parameter space, especially for the β angle which is typically much flatter than either α or γ .

5.2.3 Parameter Space Fits with Artificial Neural Networks

We now describe the architecture, training and optimization of our neural networks for the fitting coefficients of the coprocessing modes each decomposed into amplitude and phase and Euler angles, and discuss the achieved accuracy for each of component separately. We build the neural network for each model component using `Tensorflow` [194] and `Keras` [195]. Specifically, we use the `Sequential` model with fully-connected `Dense` layers. A summary of the final neural network architectures for each coprocessing mode and the Euler angles is given in Table III. As an example, a graphical representation of the neural network architecture for the coprocessing (2,2)-mode phase is shown in Fig. 5.2. The neural network is shown by the red and teal rectangles, where the red ones represent the four fully-connected hidden layers, each with 320 neurons for this component and a `Softplus` activation function, and the teal ones show the input and output layers: 4 neurons for the intrinsic parameters \mathbf{X} , and 29 for the output layer as this is the number of empirical time nodes T_i for this component. The output can then be reinterpolated onto the full uniform time grid, and inverse scaled to produce the full coprocessing (2,2)-mode phase ϕ_{22} . For this particular component we have applied scaling to the \mathbf{X} and \mathbf{Y} data, as shown by the blue rectangles. For all ANNs we use 4 input neurons, but the detailed architecture is adapted for each component.

The size of neural network differs between the coprocessing modes and Euler angles, as shown in Table III. Additionally, the coprocessing amplitudes will not undergo inverse `MinMax` scaling as we did not scale the amplitude training data in our model construction, and for the Euler angles the \mathbf{X} data is not scaled. Lastly, we note that the size of the neural network output layer will vary, as it is equal to the number of empirical time nodes for each model component.




Figure 5.2: Graphical representation of the ANN architecture for the coprocessing (2,2)-mode phase $\phi_{22}(t; \vec{\lambda})$, as an example. This neural network takes in the **Standard** scaled intrinsic binary parameters \mathbf{X} as input, and outputs the **MinMax** scaled \mathbf{Y} , a prediction of the coprocessing (2,2)-mode phase at the empirical time nodes. This output vector may then be reinterpolated onto the full uniform time grid using the empirical interpolant, and inverse **MinMax** scaled to produce the full coprocessing mode phase ϕ_{22} .

	(ℓ, m) - Amplitude & Phase	Euler Angles
X-data conditioning	Standard	None
Y-data conditioning	None (Amplitude) MinMax (Phase)	MinMax (α, γ) None (β)
Number of input neurons	4	4
Number of layers	4	9
Neurons per layer	320	128
Optimiser	Adam	AdaMax
Activation function	Softplus	Softplus
Mini-batch size	64	512
Number of training epochs	10,000	5000

Table 5.4: Details of the final neural network architecture for each component.




Figure 5.3: Training and validation losses for the $(2, 2)$ - (left) and $(2, 1)$ -mode (right), for both amplitude and phase. The loss shown is the mean squared error (MSE) as a function of training epochs. Also shown in both panels is the learning rate (red), which changes as a function of epoch as defined by Eq. (5.7).

Coprocessing Modes

When training our coprocessing mode neural networks, we use a mean squared error (MSE) loss function. This quantity can be computed over either the training dataset to monitor training progress, or the validation dataset as a control to check for over- or under-fitting. It is defined as

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N \left| \vec{y}_i^{\text{true}} - \vec{y}_i^{\text{pred}} \right|^2, \quad (5.6)$$

where \vec{y}_i^{pred} is the output from the neural network, \vec{y}_i^{true} is the true SEOBNRv4PHM data, scaled accordingly if required and at the appropriate empirical time nodes, and N is the number of points in either the training or validation set.

The final neural network architecture for the coprocessing modes is determined through optimisation via the hyperparameter sampling package `Hyperopt` [196]. We parse choices for neural network hyperparameters, as well as a maximum number of neural network evaluations. The package then creates and trains neural networks with hyperparameters drawn from these choices, and returns the best performing hyperparameters based on a metric, which we specify to be the validation loss upon completion of training. The choices for optimisable hyperparameters are as follows: activation function (Relu [197, 198], Elu [199], Tanh, Softmax, Softplus, Softsign [200]); optimiser (Adam [201], Nadam [202], Adamax [201], Adadelta [203]), number of training epochs (1000, 2000, 10000); and mini-batch size (32, 64, 128). We refer the reader to [204] for a systematic overview of activation functions and [205] for an overview of gradient descent optimization algorithms. We also compare three sizes of neural network: 4 layers with 320 neurons per layer; 9 layers with 128 neurons per layer; and 4 layers where the number of neurons per layer is the next largest power of two from the reduced basis size. We find slightly improved performance with shallower, wider architectures, and so use the 4 layer, 320 neurons per layer architecture for the final networks. We also do not use dropout in our final configurations as we find this can create a lack of stability in training leading to higher mismatches. Our final optimal neural network architecture is detailed in Table III.

Additionally, we use an adaptive learning rate as in Ref. [153] in order to achieve faster convergence and prevent overshooting of the optimal trained weights. Our learning rate takes the form

$$\tau_i = (\tau_{\text{init}} - \tau_{\text{final}})/(1 + R[i/\Delta i]) + \tau_{\text{final}}, \quad (5.7)$$

where τ_i is the learning rate at epoch i , the initial learning rate $\tau_{\text{init}} = 10^{-3}$, the final learning rate $\tau_{\text{final}} = 10^{-5}$, the decay rate $R = 10$, and our training epoch interval $\Delta i = 2000$. Thus our learning rate exhibits step-wise changes, decreasing every 2,000 epochs.

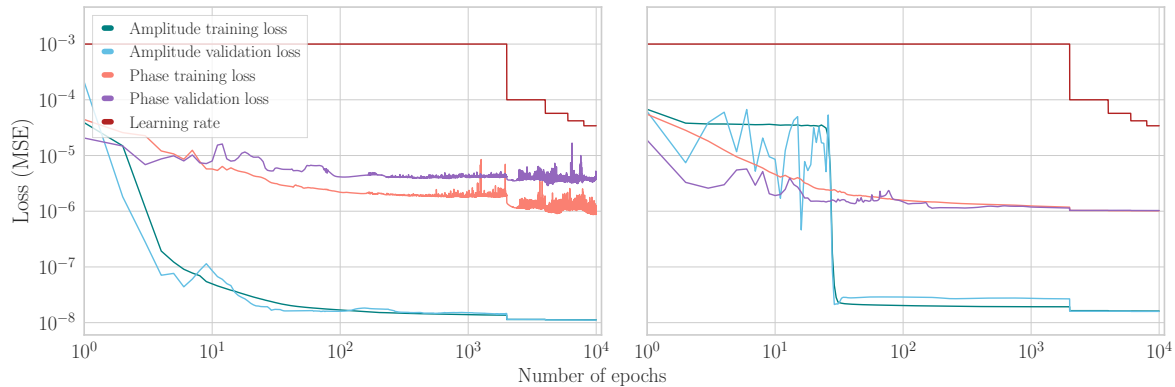


Figure 5.4: Training and validation losses for the (3, 3)- (left) and (4, 4) (right) -modes, both amplitude and phase. The loss shown is the mean squared error (MSE) as a function of epochs trained. Also shown in both panels is the learning rate, which changes as a function of epoch as defined by Eq. (5.7).

We use a mean squared error (MSE) as the loss metric, losses on both the training and validation datasets for the (2, 2) and (2, 1)-mode amplitudes and phases are shown in Fig. 5.3, as well as the variable learning rate. We see that for the (2, 2)-mode (left panel), both the amplitude and phase loss plateau around 10^{-6} after ~ 100 epochs of training, and for the (2, 1)-mode (right panel) the phase reaches a similar loss plateau as the (2, 2)-mode, however, the amplitude continues to improve to a loss value of $\sim 10^{-8}$. We also note that for all components, the training and validation losses are very comparable: a sign that we are not over-fitting in our training procedure, which would cause the validation loss to increase as the training loss continues to decrease with more training epochs. We note that in order to assess under-fitting, we would need to train the network until the training and validation loss reach a global minimum, which could take many more training epochs than shown here. The training and validation losses for the (3, 3)- and (4, 4)-modes are shown in Fig. 5.4, and we note that similarly to the (2, 1)-mode amplitude as seen in Fig. 5.3, the amplitude losses for the (3, 3)- and (4, 4)-modes evolve to a minimum of around 10^{-8} at the end of training, and the phases to around 10^{-6} .

Euler Angles

In contrast to the architecture used for the coprocessing modes, for the Euler angles we use a network that is narrower and deeper consisting of 9 layers with 128 neurons per layer. We found that the Softplus activation function coupled with the Adamax optimizer produced robust results at the desired level of accuracy, though we did not perform the more exhaustive hyperparameter optimization used in the construction of the networks for the coprocessing modes. The networks are trained for 5000 epochs using a mean squared error loss function, as defined in Eq. (5.6). For the learning rate, we use an initial value of 10^{-2} and use an adaptive scheme that reduces the learning rate when the loss has stopped improving, as implemented by the `ReduceLROnPlateau` callback in `Keras`. We found no significant improvement when exploring the use of dropout regularization or L^p regularizers³, so do not include them in the final model.

In addition to the default network above, we also constructed a neural network for the residuals between the input empirical interpolation coefficients and the default neural network predictions (see also [154]), $\tilde{y}_k = y_k^{\text{true}} - y_k^{\text{pred}}$. This allows us to reconstruct the empirical interpolation coefficients using a two step procedure: we first evaluate the default neural network then we correct for any residual errors using the second network. However, we found this gave no noticeable improvement in accuracy. Due to the additional computational cost associated to the network evaluation, we opt not to use the residuals approach in the final model.

5.2.4 Complete Surrogate Model

Once the reduced bases and empirical interpolants are built and the neural networks have been trained, we have a total of 11 surrogate models for the different components that constitute the complete precessing model, `SEOBNN_v4PHM_4dq2`: The four coprocessing modes split into amplitudes and phases, and the three Euler angles. In Fig. 5.5

³The L^p norm is defined by $\|L\|_p = \sum_n (|x_n|^p)^{1/p}$ and we applied the regularization penalty to both the kernel and bias using the `L1L2 Class` in `Keras`.

we show an example for a fiducial binary with parameters $\vec{\lambda} = \{q, \chi_{1x}, \chi_{1y}, \chi_{1z}\} = \{1.86, 0.045, -0.283, 0.274\}$, i.e. a moderately precessing binary with a moderate unequal mass ratio. We note that this particular binary was not in our training or validation datasets. The top left panel shows the mode amplitudes as predicted by the surrogate for each coprecessing mode, the top right panel the corresponding phases. The `SEOBNRv4PHM` data are shown by the dashed curves in all panels. The middle panel shows the final surrogate models for the Euler angles. We note the excellent agreement between the true data and predictions, including around merger at $t = 0M$. In the bottom panel we show the time-domain strain (Eq. (3.3)) obtained by combining the surrogate models (plus the conjugate modes) following the description in Eq. (3.12). We note, however, that we do not explicitly model the relative phase offsets between the coprecessing modes, which were incorporated manually from the true `SEOBNRv4PHM` data in the construction of the precessing strain. We leave the modelling of these relative phase offsets to future work.

Having seen the excellent agreement between prediction and true `SEOBNRv4PHM` data for a single fiducial binary, we now quantify the accuracy the surrogate models for each component across the parameter space.

For each coprecessing mode we compute white noise frequency-domain mismatches $\bar{\mathcal{M}}_f$ between the true `SEOBNRv4PHM` coprecessing waveform modes and the surrogate predictions for the test dataset, which consists of 10^4 waveforms that were not part of our training space \mathcal{T}_M (see Sec. 5.2.1). We limit our mismatch integration to start at $f_{\min} = 20$ Hz and fix the total mass to $44M_{\odot}$, which completely covers also the longest waveforms in our test set. The mismatch result for each of the four coprecessing modes is shown in in Fig. 5.6. For each of the four coprecessing modes we find that the bulk of mismatches is less than 10^{-2} or 1%, with 4.6% greater than this value for the (2, 1)-mode, 0.8% for the (3, 3)-mode, and 2.6% for the (4, 4)-mode. For the (2, 2)-mode we find that it is less than 10^{-3} with only 3.3% of mismatches greater than this, with a median mismatch of $\sim 3 \times 10^{-4}$. We find comparable performance for each of the three higher modes considered, with a median mismatch of $\sim 10^{-3}$, however we do note that there are

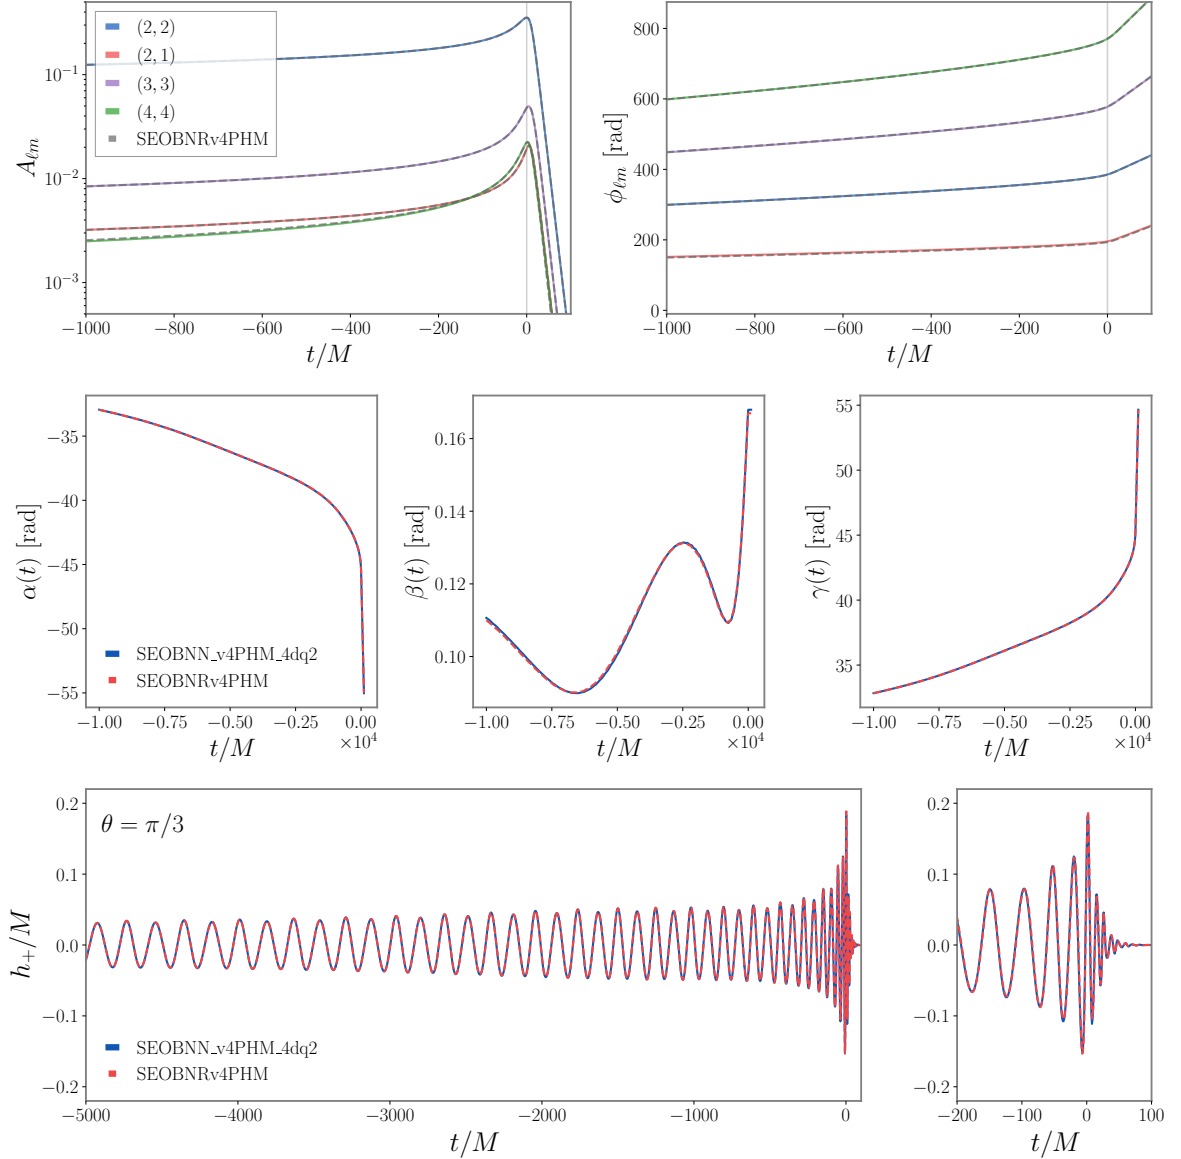


Figure 5.5: Top panel: Comparison of the coprecessing mode amplitudes (left) and phases (right) predicted by the surrogates (solid coloured lines) against the SEOBNRv4PHM data (dashed grey) for a fiducial binary with parameters $\{q, \chi_{1x}, \chi_{1y}, \chi_{1z}\} = \{1.86, 0.045, -0.283, 0.274\}$, where the Cartesian spin parameters are specified at a (2, 2)-mode reference frequency of 4 Hz. The merger at $t = 0M$ is indicated by the grey vertical line. Middle panel: Comparison of the Euler angles predicted by the neural network (blue) against the SEOBNRv4PHM data (red) for the fiducial binary. Bottom panel: The time-domain strain in the J -frame for our fiducial binary at an inclination of $\theta = \pi/3$. We include all modes up to $\ell \leq 4$.




Figure 5.6: White noise mismatches between the `SEOBNRv4PHM`-generated coprecessing frame mode data and the neural network-predicted coprecessing mode, for each of the four modes across the 10,000 binary test set. Mismatch calculations start from an $f_{\min} = 20\text{Hz}$ for a total mass of $44M_{\odot}$. Also shown for each coprecessing mode are the median mismatch (black) and 90% intervals (black dashed).

tails of higher mismatches in the odd m -modes.

Histograms of the mismatches for the coprecessing (2, 2) and (2, 1)-modes at different total masses can be found in Fig. 5.7. We choose a representative sample of four total masses $M_{\text{tot}} \in \{44, 65, 85, 125\}M_{\odot}$ and recompute mismatches in the same way as shown in Fig. 5.6, from a low frequency cutoff of 20 Hz each time, over the 10,000 binary test set, for each of the (2, 2)- (left) and (2, 1)-modes. We find that the change in total mass makes little difference in both cases, and that in fact a higher total mass than $44M_{\odot}$ slightly improves the (2, 1)-mode mismatches, as may be expected since the higher total mass is effectively a decrease in waveform length. Therefore, the results in Fig. 5.6 represents the worst case scenario.

To see where in parameter space the worst mismatches lie, particularly the high mismatch tails in the odd m -modes, we take the worst 5% for each coprecessing mode and plot them in the space of mass ratio q against $\chi_{1\text{LL}}$, with the (2, 2)- and (2, 1)-modes shown in Fig. 5.8. We see that for the (2, 2)-mode, the highest mismatches lie broadly evenly across the parameter space, although with fewer high mismatches at low in-plane spin




Figure 5.7: White noise mismatches between the `SEOBNRv4PHM`-generated coprecessing frame mode data and the neural network-predicted coprecessing mode, for the (2, 2) (left) and the (2, 1)-mode (right), as computed in Fig. 5.6 but for a range of different total masses $M_{\text{tot}} = [44, 65, 85, 125] M_{\odot}$. Mismatch calculations start from an $f_{\text{min}} = 20\text{Hz}$ and are computed across the 10, 000 binary test set.

values. For the odd (2, 1)-mode, however, the worst mismatches lie close to equal mass and at low in-plane spin values. In this region of parameter space, we expect the odd m -modes to be heavily suppressed, and so training data may be considerably more noisy, therefore leading to worse mismatches. It also means that when combining the modes into a full precessing strain, the contribution of these modes to the full signal is diminished and so will not have as much impact on the accuracy of the full waveform.

We also show in Fig. 5.9 the worst 5% of mismatches $\bar{\mathcal{M}}_f$ for the (3, 3)- (left) and (4, 4)- (right) modes. We see that for the (3, 3)-mode, similarly to the (2, 1)-mode, the worst mismatches appear around equal mass and less in-plane spin where there is less asymmetry in the system and so these modes are heavily suppressed in the full precessing strain. In contrast, and similarly to the (2, 2)-mode, the (4, 4)-mode exhibits lower mismatches overall, and more evenly spread across the parameter space, although the worst mismatches tend to be at more unequal mass ratios and larger in-plane spins.

For the Euler angles, we use time-domain mismatches, see Eq. (2.78), as the main metric to quantify the accuracy of the surrogate prediction. We show the mismatches between the `SEOBNRv4PHM` data and the surrogate models for the Euler angles in Fig. 5.10. We also demonstrate that the accuracy of the residual surrogate model outlined in Sec. 5.2.3

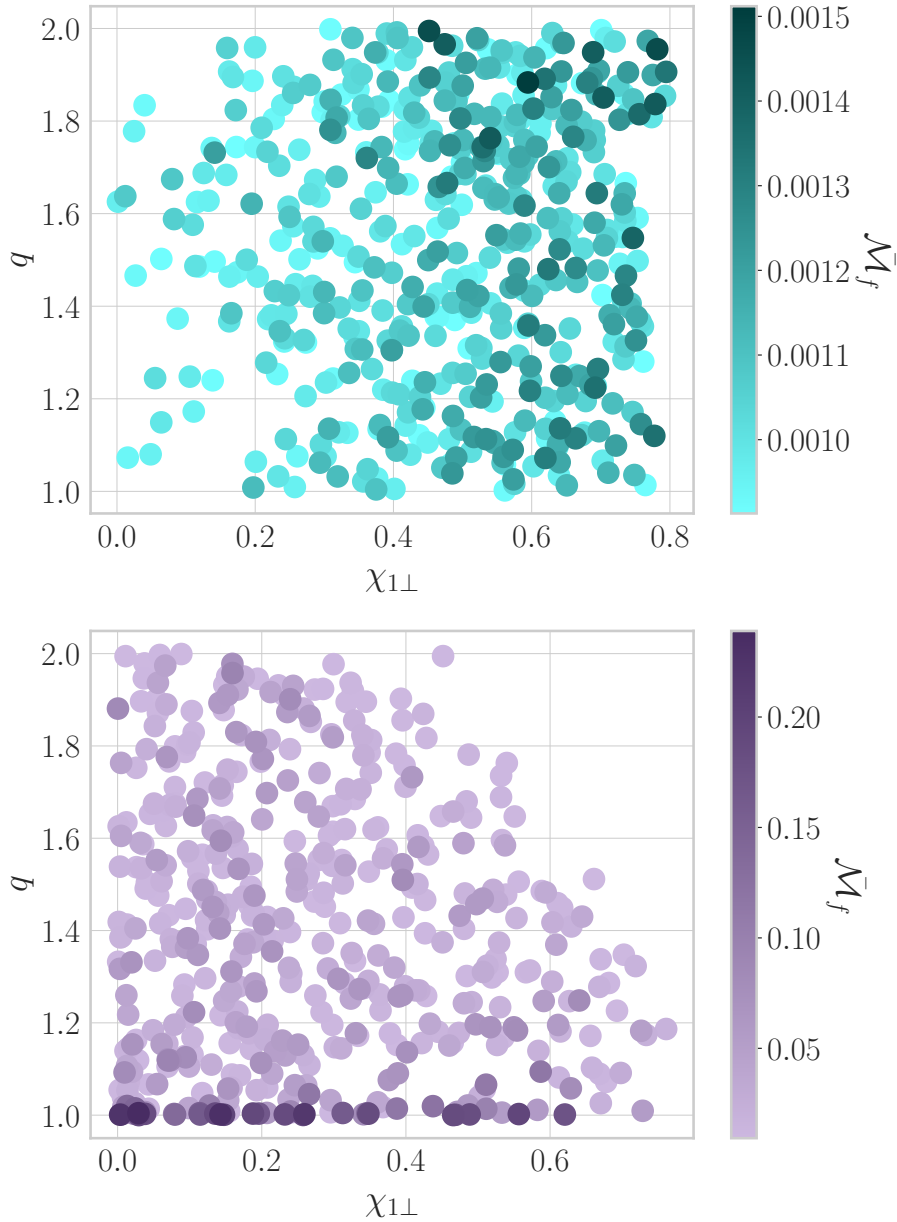


Figure 5.8: Worst 5% of test dataset mismatches for the coprecessing (2,2)- (top), (2,1)-modes (bottom), shown in parameter space of mass ratio q against in-plane spin magnitude $|\chi_{1\perp}|$. We note that the highest mismatches for the (2,2)-mode are scattered across much of this space, although with the worst mismatches at larger mass ratio and spin magnitude. In contrast, the worst mismatches for the (2,1)-mode lie in the region close to equal-mass where there is less asymmetry in the system and so this particular mode is heavily suppressed.




Figure 5.9: Worst 5% of test dataset mismatches $\bar{\mathcal{M}}_f$ for the (3, 3)- (left) and (4, 4) (right), modes, shown in parameter space of mass ratio and in-plane spin magnitude, and coloured by mismatch. We note that the highest mismatches for the (4, 4)-mode are scattered across much of this space, although with the worst mismatches at larger mass ratio and spin magnitude. In contrast, the worst mismatches for the (3, 4)-mode lie in the region close to equal-mass where there is less asymmetry in the system and so this particular mode is heavily suppressed.

offers no noticeable benefit with mismatches in broad agreement with our default model.

Finally, whilst we find it convenient to work with the $SO(3)$ representation of the Euler angles, an appealing alternative approach is to parameterize the rotation group by a set of unit quaternions [125, 206]. Fundamentally, the quaternions still describe the time-dependent rotation of the frame but are endowed with a number of beneficial mathematical properties, such as the singularities that can occur in the Euler angle formalism. For the reduced parameter space considered here, we found no noticeable benefit to adopting the quaternion framework and opted to work with Euler angles out of simplicity. We leave a more detailed investigation of the wider parameter space to future work.

5.3 Model Evaluation

5.3.1 Waveform Accuracy

The observed GW signal from single-spin precessing binary black holes depends on 12 parameters: the component masses m_i , the dimensionless spin $\vec{\chi}_1(t)$, the direction from the source frame to the observer (ι, ϕ_0) , the polarization ψ_0 , time of arrival t , the luminosity




Figure 5.10: Time-domain mismatches for the surrogate model for the Euler angles against the training data. For α , we show mismatches for the surrogate when using empirical coefficients predicted by the default network and when predicted by a two stage network that includes a fit to the residuals from the default network. We find no notable improvement in fitting the residuals.

distance d_L and the sky location (θ, ϕ) . Here we neglect the sky location and write the real-valued detector response $h_r(t)$ as

$$h_r(t) = h_+(t) \cos(2\psi) + h_\times(t) \sin(2\psi). \quad (5.8)$$

where $h(t) = h_+(t) - ih_\times(t)$. We are now interested in validating the accuracy of our surrogate model, `SEOBNN_v4PHM_4dq2`, against the slow waveform model `SEOBNRv4PHM`. To do so, we calculate strain mismatches optimized over $\{\psi, \varphi, t\}$, as these quantities are not astrophysically relevant. We follow the approaches detailed in [48, 168, 174, 192] and numerically optimize over the phase ϕ and analytically maximize over the template polarization ψ and relative time shift t ,

$$\mathcal{M}_\kappa(\phi_0^s, \psi_0^s) = \max_{t_0^h, \phi_0^h, \psi_0^h} \langle \hat{h}_r, \hat{s}_r(\phi_0^s, \psi_0^s) \rangle_f, \quad (5.9)$$

where h_r denotes the template waveform, generated by our `SEOBNN_v4PHM_4dq2` surrogate, and s_r is the signal waveform, taken to be `SEOBNRv4PHM`. We use the index κ to distinguish the match optimised over the polarisation angle from Eq. (2.74). Finally, we average the match by weighting each waveform (indexed by i) by its optimal signal-to-noise ratio ρ to account for the likelihood that the signal would have been detected. This allows us to define an orientation-averaged match as [192] (see Eq. (4.16))

$$\mathcal{M}_w = \left(\frac{\sum_i \mathcal{M}_{\kappa,i}^3 \rho_i^3}{\sum_i \rho_i^3} \right)^{1/3}, \quad (5.10)$$

and the concomitant orientation-averaged mismatch $\bar{\mathcal{M}}_w = 1 - \mathcal{M}_w$. For the match calculation, we assume a lower cutoff frequency of 20Hz and use the projected PSD for Advanced LIGO in the upcoming fourth observing run (O4) [207], consisting of the Advanced LIGO and Virgo detectors as well as KAGRA. We take the masses to be uniformly distributed between $50M_\odot$ and $200M_\odot$ and the orientation angles to be isotropic on the unit sphere. The mass ratio, spin magnitude and spin orientation are as described in Table I. We reiterate that to construct the full precessing strain from our ANN waveform model, here we use the true `SEOBNRv4PHM` relative phase offsets between the coprecessing modes. The resulting mismatches are shown in Fig. 5.11 using all $\ell \leq 4$ modes in the inertial J -frame as per Eq. (3.12). We show mismatches against the training dataset, used to construct our ANN waveform model, and the independent testing dataset to which the model has never been exposed. For both datasets we find excellent agreement and find a median mismatch of 1.9×10^{-4} . The 5th and 95th percentiles for the mismatches against the training dataset are 5.8×10^{-5} and 6.5×10^{-4} respectively. The mismatch errors here are approximately an order of magnitude below the anticipated error of `SEOBNRv4PHM` against precessing numerical relativity simulations [48]. We find that the error of our model against the input data is competitive with the accuracy provided by other surro-

gate models, e.g. [89, 160].




Figure 5.11: Orientation averaged mismatches for SEOBNRv4PHM against SEOBNN_v4PHM_4dq2 for all ≤ 4 modes in the J -frame. We show 9×10^4 binaries randomly drawn from the training (blue) dataset and 1×10^4 binaries randomly drawn from the independent testing (green) dataset, which the neural network has never been exposed to. We find excellent agreement irrespective of the dataset.

5.3.2 Timing

In order to test the efficiency of our surrogate model, we developed two interfaces. The first interface is built exclusively within the NumPy framework. The second interface uses the Tensorflow framework to provide GPU acceleration. When run on a single CPU, we find broad parity between the computational efficiency of the two implementations. However, when run on a GPU, the implementation in Tensorflow allows for significant computational speedup, as discussed below.

The typical evaluation time for a single Euler angle surrogate model is on the order of $250\mu\text{s}$. As a reminder, this includes the computational cost of producing a single prediction for the empirical interpolation coefficients from the neural network as well as the multiplication by the empirical interpolation matrix. The amplitude and phase surrogate models are marginally slower such that each waveform mode $h_{\ell m} = A_{\ell m}e^{-i\phi_{\ell m}}$

takes $\sim 925\mu\text{s}$ to generate.

In required model components are constructed from 11 individual surrogate models: 3 Euler angles and 4 waveform modes plus their conjugates. To evaluate all 3 Euler angles takes $\sim 750\mu\text{s}$ and to evaluate the 4 waveform mode surrogates takes $\sim 3.7\text{ms}$. Array conjugation is a significantly cheaper operation requiring only $\sim 10\mu\text{s}$ per array. Evaluation of the waveform modes is the single most expensive operation in our model.

Next we need to evaluate the Wigner-D matrices, $D_{mm'}^\ell(\alpha, \beta, \gamma)$, in order to perform the time-dependent rotations. This is the second most expensive operation in `SEOBNN_v4PHM_4dq2`. In order to mitigate against the computational cost, we can perform a series of optimizations, such as pre-caching of numerical coefficients. This allows us to significantly reduce the cost of evaluating the Wigner-D matrices to $\sim 5.5\text{ms}$. Further optimization could be achieved through the use of interpolating non-uniform grids or pre-compilation in C. We leave such optimizations to the future. Performing the time-dependent rotations of the waveform modes from the L -frame to the J -frame is relatively efficient, requiring only $\sim 2\text{ms}$.

Altogether, we find that the typical waveform generation cost for a signal covering the surrogate length of $10,000M$ is on the order of 18ms on a single CPU with `SEOBNN_v4PHM_4dq2`. This is on average $\mathcal{O}(10^2)$ times faster than the underlying `SEOBNRv4PHM` model, which takes $\sim 3000\text{ms}$. In addition, it is also almost three times as fast as the surrogate model presented in [208], though the surrogate model presented here is twice as long in duration spanning 10^4M compared to 5×10^3M in [208]. A notable caveat is that the surrogate model presented in [208] covers a significantly larger domain of the parameter space making any direct comparison difficult. Nevertheless, the preliminary model presented here suggests that reduced order models for precessing multipolar waveform models powered by neural networks are highly competitive relative to alternative strategies, even on a CPU. We show the typical timings for each element and for the entire waveform in Fig. 5.12. All CPU timings were generated using an Intel(R) Core(TM) i7-9750H CPU @ 2.60GHz using the NumPy interface.




Figure 5.12: Computational cost for each step in the waveform construction. A notable bottleneck is the computation of the Wigner-D matrices $D_{mm'}^\ell(\alpha, \beta, \gamma)$ over the full $10^4 M$ time grid.

However, a significant benefit of reduced order models powered by the `Tensorflow` architecture is that they provide a convenient platform for GPU acceleration. In particular, GPU acceleration is most beneficial when generating batches of surrogate models, mitigating any overhead in the transfer of data between the CPU and the GPU. Evaluating the surrogate model for the 22-mode over a varying number of binaries, we find that GPU acceleration leads to a factor $\sim \mathcal{O}(30)$ speedup in surrogate generation cost relative to CPUs. For 8192 binary configurations, we find that on a CPU each surrogate model takes ~ 30 ms compared to ~ 0.7 ms on a GPU. For the CPU-GPU benchmarking, CPU timings were performed using an Intel(R) Xeon(R) CPU @ 2.30GHz and GPU timings were performed using an NVIDIA Tesla P100-PCIE-16GB. We show the comparative CPU and GPU timings in Fig. 5.13 along with the relative speedup provided by GPU acceleration.

5.4 Discussion

As the number of observations of GW signals from BBH mergers is only set to increase with improving detector sensitivity, the availability of accurate, highly computationally efficient theoretical models is critical for future GW data analysis. Faster surrogate models of very accurate but slower underlying waveform models will prove beneficial, motivating




Figure 5.13: Computational cost per binary for evaluating the 22-mode surrogate model over batches of N_λ binaries. We show both CPU and GPU timings as well as the overall speedup enabled by GPU acceleration.

a need for exploration of novel surrogate-building methods. In this work, we have constructed a proof-of-concept time-domain surrogate model of `SEOBNRv4PHM`, which makes use of neural networks to perform parameter space fits. We follow the techniques used in Refs. [153, 209], extending them to precessing multipolar waveforms for the full inspiral-merger-ringdown signal. We decompose our surrogate model into two sections: coprecessing waveform modes defined in the non-inertial coprecessing L -frame which tracks the precessing motion of the binary, and the three Euler angles which represent the rotation between this frame and the J -frame. We consider four coprecessing modes: the dominant quadrupolar $(2, 2)$ -mode, and three higher multipoles $(2, 1)$, $(3, 3)$ and $(4, 4)$, and each of these modes is then decomposed further into amplitude and phase. Therefore we model a total of 11 components. For each component, we construct a reduced basis and empirical interpolant, before performing parameter space fits using artificial neural networks.

We demonstrate that the performance of our surrogate `SEOBNN_v4PHM_4dq2` is highly competitive in comparison to alternative surrogate modelling strategies, producing waveforms with precessing strain mismatches $\sim \mathcal{O}(10^{-3} - 10^{-4})$ against the true `SEOBNRv4PHM` data. We also show that this model is computationally efficient, producing waveforms on a CPU two orders of magnitude faster than the underlying `SEOBNRv4PHM` model, and al-

most three times as fast as the recently developed surrogate model [208] in the restricted intrinsic parameter space covered by our ANN model. We also note that our output waveforms are around two times longer than this surrogate, and that unlike the underlying `SEOBNRv4PHM` model, the evaluation time is independent upon the binary parameters. Additionally, we have shown that our surrogate model allows for an even more significant speed up in evaluation time when evaluating batches of waveforms simultaneously on GPUs.

As a proof of concept for neural network surrogates of precessing multipolar waveforms, our model is built on a restricted parameter range of mass ratios $q \in [1, 2]$ and single precessing spins $|\chi_1| \leq 0.8$, $|\chi_2| = 0$. This multidimensional portion of the precessing BBH parameter space is a starting point for surrogates which utilise neural networks, though we do not envisage any imminent roadblocks to incorporating additional information in order to extend towards the full 7D intrinsic parameter space of double precessing spins, with more unequal mass ratios. We note, however, that the size of training dataset would need to be significantly larger to accurately represents the full range of waveforms in this larger parameter space. Additionally, any higher dimensional training dataset would need to be thoroughly checked for data quality across the parameter space, as we noted that even in our restricted parameter we faced issues of pathologies in the underlying waveform model, where the coprecessing mode phases became discontinuous in the inspiral, possibly due to inaccurate next-to-quasicircular corrections in `SEOBNRv4PHM`.

To explore how accurately our model can extrapolate outside the training range, we tested each coprecessing mode surrogate on 1,500 single-spin binaries with mass ratio in $q \in (2, 4]$ or with primary spin magnitudes $0.8 < |\chi_1| \leq 0.99$ and computed mismatches against the true `SEOBNRv4PHM` data. We found that the extrapolation in spin magnitude is relatively smooth as long as the mass ratio is constrained to values that were in the original training space (i.e. $q \leq 2$), resulting in mismatches for each coprecessing mode approximately one order of magnitude worse than shown in Fig. 5.6. However, for binaries with $q > 2$, irrespective of the in-plane spin magnitude, each mode surrogate performs

poorly. The same trends were observed for the Euler angles.

Additionally, we investigated whether our model is able to capture the behaviour of binaries with two spinning black holes by using the previously developed dimensional reduction mapping of [39]. To do so, we constructed 1,000 double-spin binaries with parameters inside the training space, ensuring that the mapped spin magnitude was ≤ 0.8 . We found that the coprecessing (2,2)-mode is replicated with a mismatch accuracy of $\mathcal{O}(10^2 - 10^3)$, but that higher modes are less well reproduced.

When building our surrogate model, we explored several options to improve the accuracy of the coprecessing modes neural network fit. Before training the artificial neural networks for the coprecessing modes, we tried using principal component analysis on the reduced basis coefficient phase training data, to identify trend directions in the data which may be easier for the neural network to fit. Whilst this provided a small improvement in resulting mismatches for the (2,2)-mode, it led to marginally worse results for the (2,1)-mode and no noticeable difference in the (3,3) and (4,4)-modes. We also attempted to improve the mismatches of our coprecessing modes by training the neural networks for longer than 10,000 epochs. However, between 10,000 and 100,000 epochs, almost no improvements were seen in the loss values for both amplitudes and phases for all modes. Furthermore we tried training on the residual coprecessing phase, where the geometric mean has been subtracted to de-trend the phase data. We find this had no impact on either the reduced basis sizes or the accuracy with which we were able to train our artificial neural networks. For the Euler angles, we explored the possibility of training an additional neural network to model the residual error on the predicted α , but found no noticeable improvement.

In addition, we also explored the effect of different sizes of training data sets upon the accuracy of the coprecessing mode fits. We found that the reduced basis size and projection errors were insensitive to smaller training set sizes for sets above 100 waveforms, and similarly that the coprecessing mode mismatches for the (2,1)-mode shown in Fig. 5.6 were virtually identical when the (2,1)-phase was reconstructed on a random

training subset of 10,000. This suggests that our choice of training set size may have been conservative, and future models over this parameter space could attain similar accuracies with smaller training set sizes.

We have demonstrated the feasibility and efficacy of using neural networks as part of precessing multipolar IMR waveform surrogate models, and leave the extension to the full 7D precessing parameter space as well as the modelling of the spin evolution to further work. We suggest that with even further consideration given to neural network optimisation and data de-trending over the full 7D parameter space of generically precessing BBHs, this could prove a promising pathway towards accurate, efficient gravitational waveform surrogate model building.

6. Constraining Precession in the Black Hole Binary Population with Next-Generation Ground-Based Gravitational Wave Detectors

This chapter is a reformatted version of a paper in preparation, with the working title “Constraining Spin Precession in the Black Hole Binary Population with Next-Generation Ground-Based Gravitational Wave Detectors”. L. M. Thomas will be the lead author, and proposed the main idea behind this work, namely to assess the impact of the low frequency sensitivity on the measurement of precession in a BBH population. The details of the project were devised in discussions between L. M. Thomas, P. Schmidt and G. Pratten. The idea to perform a systematic injection series was proposed by P. Schmidt, and L. M. Thomas carried out the computational runs as detailed in Sec. 6.2, and produced all the plots. Code to produce the posterior predictive distributions was provided by G. Pratten, but L. M. Thomas generated the binary population as described in Subsec. 6.3.1, carried out all of the individual binary parameter estimation runs presented in Subsec. 6.3.2, as well as the population inference analyses presented in Subsec. 6.3.3. For the latter, L. M. Thomas implemented modifications to the Python library `gwpopulation`. L. M. Thomas, P. Schmidt and G. Pratten jointly interpreted the results presented Secs. 6.2 and 6.3. The majority of the text presented here was written by L. M. Thomas, with inputs from P. Schmidt and G. Pratten.

The confident detection of general-relativistic spin precession in a black hole binary merger with gravitational waves has so far proved elusive. While there are events such as GW190521 [37] and GW200129 [38] which show hints of possible spin precession, a unambiguous detection of precession has not yet been made. GW190521 is obfuscated by its extremely short duration signal, which leaves room for many possibilities of al-

ternative interpretations to a quasicircular precessing BBH merger [210–216], while the interpretation of GW200129 is heavily dependent on data quality and possible waveform systematics [217, 218]. However, the population of BBHs inferred after the third observing run of Advanced LIGO and Virgo hints at the presence of precession on the population level [34, 35], implying that individual precessing binaries exist in the astrophysical population. Future ground-based GW instruments with improved low-frequency sensitivity will allow us to measure the individual black hole spin magnitudes and their orientation more accurately as they will yield larger SNRs, and many more precession cycles observable in the early inspiral. Presence of precession in a BBH system allows for breaking parameter degeneracies as described in Sec. 2.4.2, and so more observed precession cycles will result in more degeneracy breaking and accurate measurements. BH spins are important tracers of a binary’s formation history, and more precise measurements will allow us to probe and constrain BBH formation channels [21–29, 31–36]. Because BBH masses are redshifted in the frame of the detector as described in Sec. 2.2.2, the improved low-frequency sensitivity of next-generation detectors will also allow us to observe massive mergers at much higher redshifts than current observations, up to $z \sim 20$ [219]. This will allow us to probe early universe BH formation and cosmology, and reveal if and how primordial black holes formed [220]. Additionally, it will improve the detectability of intrinsically high mass binaries, including intermediate-mass black hole (IMBH) binaries [221, 222], with masses of around $10^2 - 10^4 M_\odot$ as their merger frequency becomes increasingly lower. More observations of these heavy binaries will also allow us to constrain the population of hierarchical binary mergers [223], whose component BHs are the product of previous BH mergers.

The strain sensitivity of ground-based GW detectors at frequencies below 20 Hz is limited by seismic noise, gravity gradient noise, and thermal suspension noise, which are difficult to overcome. We refer the reader to Ref. [50] for a more thorough discussion of noise sources in current interferometers, as we only discuss them briefly here. While seismic noise may be limited using complex seismic isolation systems, reducing gravity




Figure 6.1: PSDs for proposed next generation detectors CE in orange, the Einstein Telescope in blue, as well as the proposed LIGO A+ upgrade in teal, and the Virgo design sensitivity curve in purple. The dashed grey line shows the location of 20Hz, which is the default low frequency cutoff for current gravitational wave analyses. At this frequency, the next generation detectors CE and ET are approximately ten times more sensitive than LIGO A+ and Virgo. The PSDs shown are from Ref. [52] for CE, Ref. [53] for ET (specifically ET-D), and Ref. [51] for LIGO A+ and Virgo design sensitivity respectively.

gradient noise may only be guaranteed with a low initial seismic excitation [224]. Therefore, choosing an interferometer site with low seismic activity is crucial to combat these low-frequency noise sources. Thermal noise in the test mass suspensions of the interferometer are the dominant thermal noise source in the frequency range 1 – 10Hz [224, 225], and may be minimised with cryogenic temperatures or different materials. The next generation of proposed GW observatories such as Cosmic Explorer (CE) [225] and the Einstein Telescope (ET) [226] are designed to have a broadband sensitivity of about a factor of ten better than Advanced LIGO A+ [225] with a much improved sensitivity below 20Hz as shown in Fig. 6.1.

However, the exact minimum frequency cutoff for planned future instruments is still uncertain. Here, we investigate the impact of the low-frequency cutoff of future GW instruments on the measurement of precession in individual as well as a population of BBHs. We describe the methodology of hierarchical Bayesian inference in Sec. 6.1, be-

fore outlining our two population spin models in Subsecs. 6.1.1 and 6.1.2. In order to disentangle whether any improvement in the spin precession measurement with a lower cutoff frequency f_{low} is due to an increase in SNR due to the increased duration of the signal, or whether it can be attributed to the increased number of observed precession cycles, we carry out a systematic study for a GW190521-like system as detailed in Subsec. 6.2. We then carry out a population study to determine the effect of changing f_{low} upon a population measurement of precession. We perform an injection and recovery of a simulated population of BBH mergers as described in Subsec. 6.3 at three different values of f_{low} . This is to assess whether the binary parameters are better recovered with increased frequency content, particularly those parameters pertaining to precession. The injected parameters of binaries used as part of this injected population are detailed in Subsec. 6.3.1. We note that the longest duration signal we analyse is around 200 seconds (corresponding to binary 11 in Tab. 6.2, from a low frequency of 5Hz), and so we do not include time- or frequency-dependent antenna patterns [227] as part of our analysis. We then use the results of the individual event analyses, discussed in Subsec. 6.3.2 to perform population inference as described in Subsec. 6.3.3, to determine whether a lower f_{low} corresponds to a better measurement of precession on a population level. Finally, we summarise our results in Sec. 6.4, discussing caveats and future work.

6.1 Hierarchical Bayesian Analysis

In this section, we give a brief overview of hierarchical Bayesian inference [228–232] to infer information about the underlying astrophysical populations of BHs, before describing the specific population hypermodels we will use in this work in Subsecs. 6.1.1 and 6.1.2. We closely follow [82] in this first part.

The goal of a population analysis is to use results from Bayesian inference on individual BBH merger events to infer properties about the population of BBH mergers as a whole. We use $\vec{\theta}$ to denote the set of 15 parameters which describe a single quasicircular BBH

merger event. We assume that for each observed event its parameters are drawn from distributions which are determined by the properties of the entire population. The set of parameters which characterises the population that the individual observations originated from, referred to as population hyperparameters, is denoted by $\vec{\Lambda}$. We also use the prefix hyper- to distinguish population quantities such as priors, posteriors and models from their individual event counterparts. The consideration of an underlying population distribution that all BBH events are drawn from is equivalent to asserting that the prior on the binary parameters $\vec{\theta}$ is conditional on $\vec{\Lambda}$:

$$\pi(\vec{\theta}) \rightarrow \pi(\vec{\theta}|\vec{\Lambda}). \quad (6.1)$$

Therefore, we wish to obtain posterior distributions for the population hyperparameters $\vec{\Lambda}$ given the GW data d . Marginalised likelihoods are obtained by integrating over the BBH parameters $\vec{\theta}$,

$$\mathcal{L}(d|\vec{\Lambda}) = \int d\vec{\theta} \mathcal{L}(d|\vec{\theta})\pi(\vec{\theta}|\vec{\Lambda}), \quad (6.2)$$

where $\mathcal{L}(d|\vec{\theta})$ is the likelihood of observing the data d given a set of BBH parameters $\vec{\theta}$. Using Bayes' theorem (see Sec.2.5), the hyperposterior for $\vec{\Lambda}$ is then given by

$$p(\vec{\Lambda}|d) = \frac{\mathcal{L}(d|\vec{\Lambda})\pi(\vec{\Lambda})}{\int d\vec{\Lambda} \mathcal{L}(d|\vec{\Lambda}) \pi(\vec{\Lambda})}, \quad (6.3)$$

where the hyperprior $\pi(\vec{\Lambda})$ reflects any prior knowledge or belief about the hyperparameters $\vec{\Lambda}$ or about the underlying astrophysical population.

GW population analysis generally involves a set of N independent BBH events, each of which has measured parameters $\vec{\theta}_i$ and associated detector data d_i . The total likelihood of observing N events can then be expressed as the product of individual event likelihoods

$$\mathcal{L}(\{d_i\}|\{\vec{\theta}_i\}) = \prod_{i=1}^N \mathcal{L}(d_i|\vec{\theta}_i). \quad (6.4)$$

Since we assume that all events were drawn from the same population, parameterised by

the same $\vec{\Lambda}$, our marginalised likelihood Eq. (6.2) becomes

$$\mathcal{L}(\{d_i\}|\vec{\Lambda}) = \prod_{i=1}^N \int d\vec{\theta}_i \mathcal{L}(d_i|\vec{\theta}_i)\pi(\vec{\theta}_i|\vec{\Lambda}), \quad (6.5)$$

and the hyperposterior of Eq. (6.3) becomes

$$p(\vec{\Lambda}|\{d_i\}) = \frac{\mathcal{L}(\{d_i\}|\vec{\Lambda})\pi(\vec{\Lambda})}{\int d\vec{\Lambda} \mathcal{L}(\{d_i\}|\vec{\Lambda})\pi(\vec{\Lambda})}. \quad (6.6)$$

The denominator in Eq. (6.6) involves integrals over the 15 BBH parameters $\vec{\theta}_i$ for N independent events, whilst also integrating over M population hyperparameters, which makes a total of $15N + M$ parameters. This equates to an extremely large prior volume, which can be prohibitively time-consuming to sample all at once. As an alternative, we may break up this process into individual integrals for each event as follows.

The posterior samples for each individual BBH event are generated with some default prior $\vec{\Phi}$, as described in Sec. 2.5, and are given by

$$p(\vec{\theta}_i|d_i, \vec{\Phi}) = \frac{\mathcal{L}(d_i|\vec{\theta}_i)\pi(\vec{\theta}_i|\vec{\Phi})}{\int d\vec{\theta}_i \mathcal{L}(d_i|\vec{\theta}_i)\pi(\vec{\theta}_i|\vec{\Phi})}, \quad (6.7)$$

which rearranges to

$$\mathcal{L}(d_i|\vec{\theta}_i) = \left(\int d\vec{\theta}_i \mathcal{L}(d_i|\vec{\theta}_i)\pi(\vec{\theta}_i|\vec{\Phi}) \right) \frac{p(\vec{\theta}_i|d_i, \vec{\Phi})}{\pi(\vec{\theta}_i|\vec{\Phi})}. \quad (6.8)$$

For convenience, we write the evidence assuming the default prior as $\mathcal{Z}_{\vec{\Phi}}(d_i)$, where

$$\mathcal{Z}_{\vec{\Phi}}(d_i) = \int d\vec{\theta}_i \mathcal{L}(d_i|\vec{\theta}_i)\pi(\vec{\theta}_i|\vec{\Phi}), \quad (6.9)$$

and Eq. (6.8) therefore reduces to

$$\mathcal{L}(d_i|\vec{\theta}_i) = \mathcal{Z}_{\vec{\Phi}}(d_i) \frac{p(\vec{\theta}_i|d_i, \vec{\Phi})}{\pi(\vec{\theta}_i|\vec{\Phi})}. \quad (6.10)$$

We may then insert this into Eq. (6.5) to obtain

$$\mathcal{L}(\{d_i\}|\vec{\Lambda}) = \prod_{i=1}^N \int d\vec{\theta}_i p(\vec{\theta}_i|d_i, \vec{\Phi}) \mathcal{Z}_{\vec{\Phi}}(d_i) \frac{\pi(\vec{\theta}_i|\vec{\Lambda})}{\pi(\vec{\theta}_i|\vec{\Phi})}. \quad (6.11)$$

We may now use the fact that when integrating over a posterior distribution, we can construct an expectation value

$$\int d\vec{\theta}_i p(\vec{\theta}_i|d_i, \vec{\Phi}) f(\vec{\theta}_i) \approx \frac{1}{n_i} \sum_k^{n_i} f(\vec{\theta}_i^k), \quad (6.12)$$

for some function f . For the i -th event we assume there are n_i posterior samples for $\vec{\theta}_i$, denoted by $\vec{\theta}_i^k$, and so the sum over k is over these posterior samples. Our marginalised likelihood from Eq. (6.11) then becomes

$$\mathcal{L}(\{d_i\}|\vec{\Lambda}) = \prod_{i=1}^N \frac{\mathcal{Z}_{\vec{\Phi}}(d_i)}{n_i} \sum_k^{n_i} \frac{\pi(\vec{\theta}_i^k|\vec{\Lambda})}{\pi(\vec{\theta}_i^k|\vec{\Phi})}, \quad (6.13)$$

where the product i is over the N independent events, and the sum k is over the n_i posterior samples of $\vec{\theta}_i^k$ for event i . If we assume a sufficiently uninformative prior $\vec{\Phi}$ when performing our individual event analyses, we can recycle our posterior samples from the individual events to obtain a posterior on $\vec{\Lambda}$, as the choice of $\vec{\Phi}$ effectively cancels out. Our approach is therefore to perform parameter estimation on individual events to produce individual event posteriors, and utilise Eqs. (6.6) and (6.13), sampling over $\vec{\Lambda}$ to obtain a hyperposterior for the population distribution. We caveat that we do not make use of any selection effects in our population analysis [233–235].

All of the above discussion has implicitly assumed a functional form of the population hyperprior in the analysis, as in order to parameterise a distribution using hyperparameters, one needs to know what form that distribution takes. As with single event parameter estimation, this functional form is a choice which we call the hypermodel. For our inference, we make use of two distinct hypermodels to infer information about the distribution of spin parameters across the binary population. We use parameterised models to con-

strain the distribution of individual spin magnitudes, and also the distribution of effective spin parameters $\{\chi_{\text{eff}}, \chi_p\}$. In Subsecs. 6.1.1 and 6.1.2 we discuss each of these hypermodels in more detail.

6.1.1 Hypermodel of Spin Magnitude Distribution

Our first hypermodel uses a phenomenological approach to model the distribution of spin magnitudes χ_1 and χ_2 . This hypermodel was first introduced in Ref. [32] and subsequently used in Refs. [33–35] in conjunction with a hypermodel for the spin tilts introduced in Ref. [29]. We note that we will focus exclusively on the spin magnitudes hypermodel in this work due to small number statistics, as will be discussed in Sec. 6.3.1, and leave the spin tilts analysis to future work.

This spin magnitudes hypermodel, which we call the **Beta Magnitudes** hypermodel, assumes the distribution of each spin magnitude χ_1, χ_2 to follow a beta distribution,

$$p(\chi_i | \alpha_{\chi_i}, \beta_{\chi_i}) = \frac{\chi_i^{\alpha_{\chi_i}-1} (\chi_{\text{max}} - \chi_i)^{\beta_{\chi_i}-1}}{\mathbf{B}(\alpha_{\chi_i}, \beta_{\chi_i}) \chi_{\text{max}}^{\beta_{\chi_i} + \alpha_{\chi_i} + 1}}, \quad (6.14)$$

where $i = 1, 2$, and χ_{max} is the maximum spin value which we take to be 1 in all our analysis. The quantity $\mathbf{B}(\alpha_{\chi_i}, \beta_{\chi_i})$ is a normalisation constant,

$$\mathbf{B}(\alpha_{\chi_i}, \beta_{\chi_i}) = \frac{\Gamma(\alpha_{\chi_i})\Gamma(\beta_{\chi_i})}{\Gamma(\alpha_{\chi_i} + \beta_{\chi_i})}, \quad (6.15)$$

where Γ is the gamma function, and $\alpha_{\chi_i}, \beta_{\chi_i} > 0$. Therefore this distribution depends on four hyperparameters, the shape parameters α_1, β_1 for the primary spin magnitude distribution and α_2, β_2 for the secondary spin magnitude distribution.

6.1.2 Hypermodel of Effective Spin Parameter Distribution

Our second spin hypermodel assumes the distribution of the effective spin parameters $\chi_{\text{eff}}, \chi_p$ to be a bivariate Gaussian G [35]. This hypermodel, based on an earlier iteration

in Ref. [236] and subsequently adapted for use in Refs. [34, 35], makes use of the fact that the effective spins are often better constrained in individual event analyses than the individual BH spin magnitudes or tilts. The joint probability density function is given by

$$p(\chi_{\text{eff}}, \chi_p | \mu_{\chi_{\text{eff}}}, \sigma_{\chi_{\text{eff}}}, \mu_{\chi_p}, \sigma_{\chi_p}, \rho) \propto G(\chi_{\text{eff}}, \chi_p | \vec{\mu}, \Sigma), \quad (6.16)$$

where $(\mu_{\chi_{\text{eff}}}, \mu_{\chi_p})$ are the means, and Σ is the covariance matrix given by

$$\Sigma = \begin{pmatrix} \sigma_{\chi_{\text{eff}}}^2 & \rho \sigma_{\chi_{\text{eff}}} \sigma_{\chi_p} \\ \rho \sigma_{\chi_{\text{eff}}} \sigma_{\chi_p} & \sigma_{\chi_p}^2 \end{pmatrix}. \quad (6.17)$$

Therefore, this hypermodel depends on five hyperparameters: $\mu_{\chi_{\text{eff}}}$ and μ_{χ_p} which are the means of the χ_{eff} and χ_p one-dimensional marginalised distributions respectively; $\sigma_{\chi_{\text{eff}}}$ and σ_{χ_p} the standard deviations of these distributions; and ρ which parameterises the degree of correlation between the χ_{eff} and χ_p distributions. As χ_{eff} varies in the region $[-1, 1]$, $\mu_{\chi_{\text{eff}}}$ may also take these values, and similarly for χ_p and μ_{χ_p} over $[0, 1]$. The bivariate Gaussian function G is also truncated and normalized accordingly to satisfy the allowed ranges for the effective spins.

6.2 Systematic Injection Study

A key consideration when assessing possible improvement in the parameter estimation results with decreasing f_{low} is to disentangle the effects of increasing SNR and more precession cycles, which both can lead to better spin constraints. In this section, we perform a systematic series of waveform injections and recoveries to disentangle these two effects. We describe the setup and methodology of the analysis in Subsec. 6.2.1, before presenting our results in Subsec. 6.2.2.

6.2.1 Methodology

To disentangle these two effects, we perform full Bayesian inference (see Sec. 2.5) on a systematic series of injections for a binary with parameters similar to that of the recovered parameters for the observed event GW190521 [37], which showed possible posterior support for presence of precession, but due to the heavy BH masses was too short a signal to effectively constrain the binary spin parameters [37]. The injected parameters are given in Tab. 6.1. We perform two complementary series of injections, the first as shown in columns 3-5 of Tab. 6.1 where the injected luminosity distance d_L remains fixed, and so the SNR increases as the f_{low} is lowered from 20Hz, to 10Hz, and then to 5Hz. We then inject a second series (as shown in columns 6-8 of Tab. 6.1) where the luminosity distance d_L is adjusted to keep the SNR fixed across the three values of f_{low} , to isolate the effect of only increasing number of precession cycles without changing the SNR. We note that to save on computational time, we use the same analysis for $f_{\text{low}} = 5\text{Hz}$ for each of the two systematic series, so in total we only perform 5 injection and recovery analyses across the two series.

We calculate the number of precession cycles and SNR as follows: An approximate expression for the number of precession cycles between f_{low} and some maximum frequency $f_{\text{max}} < f_{\text{merger}}$, for a BBH with total mass M (in solar masses M_\odot at 2PN order is given in Eq. [45] of Ref. [14] by

$$\alpha(f_{\text{low}}, f_{\text{max}})/2\pi \approx \begin{cases} 11 \left(1 + \frac{3m_1}{4m_2}\right) \frac{10M_\odot}{M} \frac{f_{\text{low}}}{f_{\text{max}}}, & |\vec{L}| \gg |\vec{S}|, \\ 1.9 \left(1 + \frac{3m_1}{4m_2}\right) \frac{m_1}{m_2} \frac{|\vec{S}|}{m_1^2} \left(\frac{10M_\odot}{M} \frac{f_{\text{low}}}{f_{\text{max}}}\right)^{2/3}, & |\vec{S}| \gg |\vec{L}|. \end{cases} \quad (6.18)$$

We relax the conditions in the above equation to be $|\vec{L}| \geq |\vec{S}|$ and $|\vec{S}| > |\vec{L}|$, and verify them for each binary using the expression for $|\vec{L}|$ to 3PN for aligned spins as given in Ref. [148] and derived in Refs. [57, 237, 238]:

$$|\vec{L}| = \frac{\eta M^2}{v} \left[L_0 + L_1 v + L_2 v^2 + L_3 v^3 + L_4 v^4 + L_5 v^5 + L_6 v^6 + \dots \right] \quad (6.19)$$

6. CONSTRAINING PRECESSION IN THE POPULATION

	Injected value	Fixed d_L			Fixed ρ_{network}		
		5Hz	10Hz	20Hz	5Hz	10Hz	20Hz
$m_{1,\text{source}} [M_\odot]$	84.83	84.84 ^{+0.72} _{-0.72}	85.01 ^{+1.04} _{-0.80}	84.92 ^{+1.26} _{-1.23}	84.84 ^{+0.72} _{-0.72}	84.89 ^{+0.80} _{-0.75}	84.82 ^{+0.88} _{-0.91}
$m_{2,\text{source}} [M_\odot]$	65.86	65.88 ^{+0.56} _{-0.56}	65.88 ^{+0.88} _{-0.89}	65.60 ^{+1.20} _{-1.27}	65.88 ^{+0.56} _{-0.56}	65.85 ^{+0.82} _{-0.86}	65.85 ^{+0.94} _{-0.96}
χ_1	0.69	0.69 ^{+0.03} _{-0.03}	0.69 ^{+0.09} _{-0.21}	0.43 ^{+0.28} _{-0.24}	0.69 ^{+0.03} _{-0.03}	0.69 ^{+0.05} _{-0.22}	0.68 ^{+0.06} _{-0.09}
χ_2	0.73	0.72 ^{+0.05} _{-0.05}	0.72 ^{+0.11} _{-0.07}	0.67 ^{+0.27} _{-0.59}	0.72 ^{+0.05} _{-0.05}	0.72 ^{+0.09} _{-0.06}	0.71 ^{+0.07} _{-0.06}
θ_1 [deg]	81.54	81.40 ^{+2.39} _{-2.49}	81.19 ^{+3.45} _{-3.82}	75.44 ^{+8.32} _{-42.69}	81.40 ^{+2.39} _{-2.49}	81.93 ^{+4.01} _{-4.01}	81.36 ^{+4.01} _{-4.58}
θ_2 [deg]	85.37	85.40 ^{+3.99} _{-3.77}	84.84 ^{+6.20} _{-5.75}	87.95 ^{+20.50} _{-11.79}	85.40 ^{+3.99} _{-3.77}	84.80 ^{+5.16} _{-5.73}	85.94 ^{+7.45} _{-6.30}
χ_P	0.68	0.68 ^{+0.03} _{-0.03}	0.68 ^{+0.09} _{-0.10}	0.56 ^{+0.16} _{-0.31}	0.68 ^{+0.03} _{-0.03}	0.68 ^{+0.06} _{-0.07}	0.67 ^{+0.05} _{-0.07}
χ_{eff}	0.08	0.08 ^{+0.01} _{-0.01}	0.09 ^{+0.02} _{-0.01}	0.08 ^{+0.03} _{-0.04}	0.08 ^{+0.01} _{-0.01}	0.08 ^{+0.01} _{-0.01}	0.08 ^{+0.02} _{-0.03}
ϕ_{12} [deg]	190.77	191.53 ^{+6.97} _{-6.13}	195.03 ^{+29.87} _{-13.92}	219.46 ^{+76.83} _{-69.10}	191.53 ^{+6.97} _{-6.13}	193.22 ^{+32.12} _{-11.47}	194.55 ^{+20.62} _{-12.72}
α [deg]	198.10	198.20 ^{+2.47} _{-1.74}	198.16 ^{+2.35} _{-1.68}	198.10 ^{+2.20} _{-1.61}	198.20 ^{+2.47} _{-1.74}	198.15 ^{+2.13} _{-1.54}	197.98 ^{+1.60} _{-1.30}
δ [deg]	-44.83	-45.04 ^{+11.06} _{-10.01}	-44.79 ^{+9.83} _{-9.21}	44.43 ^{+9.83} _{-9.21}	-45.04 ^{+11.06} _{-10.01}	-44.77 ^{+8.80} _{-8.77}	-43.97 ^{+7.40} _{-6.47}
θ_{JN} [deg]	44.49	44.39 ^{+1.26} _{-1.31}	44.25 ^{+2.37} _{-3.51}	38.03 ^{+7.21} _{-6.36}	44.39 ^{+1.26} _{-1.31}	44.27 ^{+1.97} _{-2.61}	44.01 ^{+2.24} _{-2.65}
ψ [deg]	104.21	103.86 ^{+2.13} _{-2.22}	103.64 ^{+10.80} _{-4.21}	106.48 ^{+11.44} _{-8.20}	103.86 ^{+2.13} _{-2.22}	103.86 ^{+9.56} _{-2.91}	103.49 ^{+3.42} _{-3.72}
z		0.82	0.82	0.82	0.82	0.71	0.57
d_L [Mpc]		5296.7	5296.7	5296.7	5296.7	4468.7	3393.3
ρ_{network}		374.2	315.7	239.7	374.2	374.2	374.2
$\alpha(f_{\text{low}}, f_{\text{ISCO}})$		2.64	1.21	0.49	2.64	1.21	0.49

Table 6.1: Injected and recovered parameters for the systematic series for a GW190521-like binary, including source frame masses, spin, and extrinsic parameters. All spin and orientation-related parameters are defined at a reference frequency of $f_{\text{ref}} = 5\text{Hz}$. We perform two series of injections, firstly a set where the luminosity distance d_L is fixed and so the SNR ρ_{network} increases as f_{low} decreases. For the second set we fixed the SNR ρ_{network} by adjusting d_L , to isolate the effect of changing number of precession cycles. Columns 3-5 show the injected d_L , redshift z and ρ_{network} , and recovered parameters for the fixed d_L series, and columns 6-8 show the same for the fixed ρ_{network} series. The recovered values show the medians, and the 90% credible intervals.

where, as in Sec. 2.2.1, M is the total mass, η is the symmetric mass ratio, and v is the

velocity as in Eq. (2.83). The post-Newtonian coefficients L_i are given by

$$L_0 = 1, \quad (6.20)$$

$$L_1 = 0, \quad (6.21)$$

$$L_2 = \frac{3}{2} + \frac{\eta}{6}, \quad (6.22)$$

$$L_3 = \chi_{1\parallel} \left(-\frac{5}{3} - \frac{5\delta}{3} + \frac{5\eta}{6} \right) + \chi_{2\parallel} \left(-\frac{5}{3} + \frac{5\delta}{3} + \frac{5\eta}{6} \right), \quad (6.23)$$

$$L_4 = \frac{27}{8} - \frac{19\eta}{8} + \frac{\eta^2}{24}, \quad (6.24)$$

$$L_5 = -\frac{7}{144} \left(\chi_{1\parallel} [72 + \delta(72 - 31\eta) + \eta(-121 + 2\eta)] + \chi_{2\parallel} [72 + \eta(-121 + 2\eta) + \delta(-72 + 31\eta)] \right), \quad (6.25)$$

$$L_6 = \frac{1}{1296} (10935 + \eta(-62001 + \eta(1674 + 7\eta)) + 4428\pi), \quad (6.26)$$

where $\delta = \sqrt{1 - 4\eta}$ and $\chi_{i\parallel}$ are the components of the BH spins aligned with the orbital angular momentum (see Sec. 2.4). We then set our maximum frequency to be the frequency of the Kerr *innermost stable circular orbit* (ISCO) f_{ISCO} . Here we provide a brief outline of the methodology we use, but for more details and derivations of the expressions we refer the reader to Ref. [239]. More specifically, we choose f_{ISCO} according to the orbital frequency of a (non-spinning) test particle orbiting a Kerr black hole at the ISCO. The Kerr black we take to be the remnant of the BBH merger, with mass equal to the total mass of the binary and spin equal to the remnant spin. The f_{ISCO} is given by

$$f_{\text{ISCO}} = \left(\pi r_{\text{ISCO}}^{3/2} \chi_f \right)^{-1}, \quad (6.27)$$

where χ_f is the dimensionless spin of the merger remnant black hole. We obtain the value of χ_f using the phenomenological fits to η , $\chi_{1\parallel}$ and $\chi_{2\parallel}$ from Ref. [145], where we specify the spin components $\chi_{1\parallel}$, $\chi_{1\parallel}$ at our reference frequency of 5Hz. The ISCO radius r_{ISCO}

for a Kerr black hole is given by

$$r_{\text{ISCO}} = M \left(3 + Z_2 \mp \sqrt{(3 - Z_1)(3 + Z_1 + 2Z_2)} \right), \quad (6.28)$$

where the choice of sign is negative if the remnant spin vector $\vec{\chi}_f$ is aligned with the orbital angular momentum \vec{L} (as defined at the reference frequency), ie. the angle between them is less than 90 degrees, and positive otherwise. The quantities Z_1 and Z_2 are given by

$$Z_1 = 1 + (1 - \chi_f^2)^{1/3} \left[(1 + \chi_f)^{1/3} + (1 - \chi_f)^{1/3} \right], \quad (6.29)$$

$$Z_2 = \sqrt{3\chi_f^2 + Z_1^2}. \quad (6.30)$$

The ISCO frequency is an approximation to the final GW frequency before plunge, and so a natural endpoint to the precessional evolution of the binary. It also signifies the breakdown of the PN series. The total number of precession cycles calculated in this way for our systematic injection series is shown in Tab. 6.1.

Considering a binary with intrinsic parameters $\vec{\theta}$, and resulting waveform $h(t; \vec{\theta})$, the (optimal) SNR ρ is given by

$$\rho^2 = \langle h, h \rangle = 4\mathcal{R} \int_{f_{\text{low}}}^{f_{\text{max}}} \frac{\tilde{h}(f)\tilde{h}^*(f)}{S_n(|f|)} df. \quad (6.31)$$

As the integrand is positive definite, the SNR ρ will increase as f_{low} decreases, assuming f_{max} remains fixed. There is an expected approximate scaling between SNR and the posterior width Γ on recovered parameters [240, 241],

$$\Gamma_{90\%}(\theta) \propto \frac{1}{\sqrt{\rho}}, \quad (6.32)$$

where $\Gamma_{90\%}$ denotes the width of the 90% credible interval on the 1D marginalised posterior of a particular parameter $\theta \in \vec{\theta}$. Therefore, as the SNR increases with decreasing f_{low} , we expect the posterior widths to decrease accordingly, and hence an increase in the

precision to which binary parameters can be measured. Additionally we may also expect the posterior widths to decrease due to an increased number of observed precession cycles, particularly for the spin parameters. Presence of precession in a binary breaks parameter degeneracies as described in Sec. 2.4.2, and so the observation of more precession cycles will allow for more degeneracy breaking likely resulting in smaller posterior widths.

For the systematic analysis we use a three-detector network consisting of one triangular-shaped ET-D [224] at the current location of Virgo, and two interferometers at the location of Hanford and Livingston with LIGO A+ design sensitivity [225]. We only vary the lower cutoff frequency in the next generation detector in the network, in this case ET, while keeping the low frequency at 20Hz for the Hanford and Livingston A+ interferometers. The network SNRs ρ_{network} shown in Tab. 6.1 are the injected network SNRs. Assuming a coherent GW signal, the network SNR ρ_{network} is defined by the quadrature sum of the SNRs in each detector in the network [242]. For a network of n_D detectors, each with SNR ρ_i , this is given by

$$\rho_{\text{network}}^2 = \sum_{i=1}^{n_D} \rho_i^2, \quad (6.33)$$

and for our network of ET-D, H+ and L+ we have

$$\rho_{\text{network}}^2 = \rho_{\text{ET-D}}^2 + \rho_{\text{H+}}^2 + \rho_{\text{L+}}^2. \quad (6.34)$$

We use the parameter estimation code `Bilby` [87] with the nested sampler `Dynesty` [86] to perform our analyses. We perform zero noise injections with the `IMRPhenomXPHM` waveform model [168] and recover with the same model, to avoid incurring biases in our parameter estimation due to waveform systematics. We use uninformative priors uniform in chirp mass, mass ratio, and spin magnitudes, isotropic in spin orientation, and the standard prior for extrinsic parameters (see e.g. [134]).




Figure 6.2: Posterior distributions for the spin parameters of the GW190521-like injected binary whose parameters are shown in Tab. 6.1, for the fixed d_L systematic series. We show 90% credible 2D contours and 1D distributions for each of the 3 f_{low} values, 20Hz in purple, 10Hz in teal, and 5Hz in orange, with the true injected values shown in black. As a result of both the increasing SNR and the higher number of observed precession cycles as f_{low} decreases, the spin posterior widths/contours shrink accordingly, also breaking degeneracies and reducing multimodalities present in the analysis from 20Hz.

6.2.2 Results

We show the results for the systematic series study in Figs. 6.2 and 6.3. In Fig. 6.2, we show the posterior distributions for the spin parameters at a reference frequency of $f_{\text{ref}} = 5\text{Hz}$ in the systematic series with fixed d_L , whose parameters are detailed in columns

3-5 of Tab. 6.1. We show the 90% credible contours (see Sec. 2.5) and distributions for $f_{\text{low}} = 20\text{Hz}$ in purple, $f_{\text{low}} = 10\text{Hz}$ in teal, and $f_{\text{low}} = 5\text{Hz}$ in orange, with the true injected values shown in black. As the SNR ρ_{network} and number of observed precession cycles both increase with lower f_{low} , the width of the 90% credible intervals decreases as expected, and we have verified that the widths scale according to Eq. (6.32).

The 90% credible interval width of the 1D posteriors of χ_p shrinks from 0.48 at $f_{\text{low}} = 20\text{Hz}$ (nearly half the possible range), to 40% of this width at 10Hz (0.19), and then 13% of the 20Hz width at 5Hz (0.06). The χ_{eff} posterior widths reduce by 77% between 20Hz and 5Hz. The largest overall improvement for any spin parameter is seen in the primary spin tilt θ_1 , where the 90% credible interval shrinks from 51 degrees at 20Hz, down to 7 degrees at 10Hz and 5 degrees at 5Hz, corresponding to a 90% reduction in the posterior width. The secondary tilt θ_2 reduces from a 90% credible interval of 32 degrees at 20Hz, down to 37% of this width at 10Hz with 12 degrees, and 24% at 5Hz with 8 degrees. Finally, the spin magnitudes have a width of 0.52 and 0.86 at 20Hz for χ_1 and χ_2 respectively, both more than half the total prior width, and these reduce to 0.30 and 0.19 at 10Hz, and then to 0.06 and 0.09 at 5Hz. For both magnitudes this is around a 90% reduction in the credible interval size. We note that for the results with $f_{\text{low}} = 20\text{Hz}$ (purple), the large credible interval width are not due to poor convergence of the runs, which we verified by choosing (even) more conservative sampler settings. It is evident from Fig. 6.2 that the posteriors for several spin parameters cover almost the entire prior range, and hence we conclude that these parameters cannot be constrained effectively for those injections. This is also consistent with the majority of spin constraints obtained to date from GW observations [218]. This could be due to a shorter duration signal or lower SNR than the 10Hz and 5Hz results shown on the same plot, or a combination of these two effects. Dramatic though these improvements are, it is not yet clear whether they are as a result of the increasing SNR, increased number of precession cycles, or a combination of the two.

To definitively answer this question, we performed the second set of injections with




Figure 6.3: Posterior distributions for spin parameters of the GW190521-like injected binary whose parameters are shown in Tab. 6.1, for the series with fixed network ρ_{network} . We show 90% credible 2D contours and 1D distributions for each of the 3 f_{low} values, 20Hz in purple, 10Hz in teal, and 5Hz in orange, with the true injected values shown in black. The luminosity distances d_L have been adjusted such that each of the three injections has the same SNR $\rho_{\text{network}} \approx 374$, as detailed in Tab. 6.1. As f_{low} is lowered, the increased number of precession cycles leads to reduce the uncertainty on the recovered parameters.

fixed network SNR $\rho_{\text{network}} = 374.2$, whose parameters are detailed in columns 6-8 of Tab. 6.1. The resulting posterior distributions for the spin parameters are shown in Fig. 6.3. As before, we show the 90% credible contours and distributions for $f_{\text{low}} = 20\text{Hz}$ in purple, $f_{\text{low}} = 10\text{Hz}$ in teal, and $f_{\text{low}} = 5\text{Hz}$ in orange, with the true injected values shown in black. We find small improvements in the posterior widths as f_{low} decreases,

which is now purely due to a change in the number of precession cycles in band and not the SNR, which remains fixed. The recovered posteriors mostly show decreasing 90% credible interval widths with decreasing f_{low} . For example, the spin tilts θ_1 and θ_2 have a 90% credible interval of 9 and 14 degrees respectively with $f_{\text{low}} = 20\text{Hz}$. This reduces to 8 and 11 degrees respectively at 10Hz, and 5 and 8 degrees at 5Hz, which is an improvement of around 56% between the 20Hz and 10Hz width for each tilt angle. The effective spin χ_{eff} , which already has narrow posterior at 20Hz with a 90% width of only 0.04, reduces to 0.02 at 5Hz corresponding to a 50% reduction in uncertainty. Interestingly, the individual BH spin magnitudes χ_1 and χ_2 and the effective precession spin χ_p all shows signs of being recovered slightly worse from 10Hz than 20Hz due to a slight bimodality in the 10Hz posteriors (teal) as can be seen in Fig. 6.3. However, the 5Hz posteriors all show dramatic improvements with respect to the 20Hz widths, with the 90% credible interval on χ_p reducing by 52%, and χ_1 and χ_2 reducing by 60% and 41% respectively.

We note that the recovered 20Hz spin posteriors are now much tighter than in Fig. 6.2 and no longer uninformative. While the injected parameters for these two analyses are the same including the length of signal and number of precession cycles, except for the luminosity distance, this results in a difference in the network SNR ρ_{network} of ≈ 150 . Therefore, the reduced uncertainty on the recovered parameters in Fig. 6.3 in comparison to Fig. 6.2, shown by the decreased 90% credible intervals between columns 5 and 8 of Tab. 6.1, is the direct result of the increased SNR. However, we note that in general, as we see measurement improvements with decreasing f_{low} both with increasing SNR in Fig. 6.2, and when the SNR is kept fixed in Fig. 6.3, then any improvements with a lower frequency cutoff must be the result of a combination of higher SNR and more precession cycles observed.

With these results in mind, we will now assess the potential improvements in estimating the hyperparameters $\vec{\Lambda}$ of an astrophysical population of black holes with decreasing f_{low} , as described in Sec. 6.3.

6.3 Population Study

In Sec. 6.2, we showed that a lower minimum frequency cutoff f_{low} can lead to significant improvements in the precision of the inferred spin parameters of black hole binaries. We also showed that, while this is mainly due to the increased SNR of longer-in-band signals, it can also be attributed in part to the observation of more precession cycles. We now apply these findings to a population level analysis, to determine whether the improvements shown in the measurement of individual binary parameters translate into an improved recovery of the population parameters.

To do so, we first draw a subset of BBHs from the astrophysical population of black holes after O3 [35]. We then perform Bayesian inference on each binary using three different low cutoff frequencies, 5Hz, 10Hz, and 20Hz, similarly to the systematic series described above. We note that as the luminosity distance d_L for each binary each kept constant across the three frequencies, the SNR for each binary will vary, similarly to the series as shown in Fig. 6.2. Subsequently, following the process detailed in Sec. 6.1, we use the parameter estimation samples to perform hierarchical population analyses on the spin parameters to infer the underlying spin distribution. For each hypermodel we perform three analyses, one for each value of f_{low} . Thus, we can directly assess the impact of varying f_{low} on the inferred spin population.

For our individual event analyses we use a three-detector network consisting of one CE at the current location of Hanford, one interferometer at the location of Livingston with LIGO A+ design sensitivity (referred to as L+), and one interferometer at the location of Virgo at Advanced Virgo design sensitivity [243]. Similarly to the systematic series, we use the public inference library Bilby [87] with the nested sampler Dynesty [86] to perform our analyses, and inject and recover with the same waveform model IMRPhenomXPHM [168] into zero noise. We only vary the low frequency cutoff for the next generation detector in our network, in this case CE, while keeping f_{low} at 20Hz for the other two interferometers. For binaries with signal durations longer than four seconds, we make use of a multibanded likelihood [244] to accelerate likelihood evaluations by using frequency grids with adaptive

resolutions such that there are fewer evaluation points at lower frequencies, where the binary frequency evolution is slower (see Eq. (2.59)). If the signal duration in band is less than 4 seconds, we use the standard BBH likelihood Eq. (2.103) as implemented in `Bilby`. We use the same priors as for the systematic series analysis as described in Subsec. 6.2.1.

In the following subsections, we describe our population analysis in more detail. In Subsec. 6.3.1 we outline how we draw our subset of binaries from the current best knowledge of the astrophysical distribution of BHs. In Subsec. 6.3.2 we show a selection of results for the individual binary parameter estimation analyses. Then in Subsec. 6.3.3, we discuss how we perform the population analyses and present the results to ascertain whether the improvements in accuracy of recovered binary parameters when decreasing f_{low} leads to an improvement in how precisely we can measure the population hyperparameters.

6.3.1 Population Binary Parameters

We draw our set of individual binaries from the inferred population from GWTC-3 [35] using the hyperparameter posterior samples which were obtained in therein, assuming a `Power Law + Peak` hypermodel for the BH masses and the `Default` spin hypermodel. The `Power Law + Peak` hypermodel uses two key ingredients: The first is a power law distribution for the BH mass between some minimum and maximum values m_{min} and m_{max} , outside of which there is a sharp cutoff. The second is a Gaussian component motivated by the fact that the mass loss undergone by pulsational pair-instability supernovae could lead to a larger number of BHs with masses smaller than the pair-instability gap, which lies between $\sim 50M_{\odot}$ and $\sim 130M_{\odot}$ [245]. For more details of this hypermodel we refer the reader to Ref. [29] where the hypermodel was introduced, and Refs. [33–35] where it is referred to as ‘Model B’, ‘Model C’, and `Power Law + Peak` respectively.

The `Default` spin hypermodel is a combination of: (i) a hypermodel for the spin tilts [246], assuming a mixture of an isotropic, dynamically formed subpopulation, and a ‘field’ subpopulation whose spins are approximately aligned with the orbital angular momentum [247] (for a review of BBH formation channels we refer the reader to, eg. Refs. [21, 70]);

and (ii) the **Beta Magnitudes** hypermodel described in Sec. 6.1.1. In our analysis, we will exclusively focus on the spin magnitudes. The GWTC-3 analysis assumes that the primary and secondary BH spin magnitudes are drawn from the same Beta distribution, which enforces that $\sigma_1 = \sigma_2 = \sigma$, as well as $\alpha_{\chi_1} = \alpha_{\chi_2} = \alpha_\chi$ and $\beta_{\chi_1} = \beta_{\chi_2} = \beta_\chi$. Henceforth we refer to the combination of the **Power Law + Peak** mass hypermodel and **Default** spin hypermodel as the **GWTC-3** hypermodel. Given a value for each of the hyperparameters for the **GWTC-3** hypermodel, we obtain a probability distribution for each binary parameter in $\vec{\theta}_{\text{GWTC-3}} = \{m_1, m_2, \chi_1, \chi_2, \theta_1, \theta_2\}$, from which we draw our subset of binaries. The distributions of the population hyperparameters, which include a set of hyperparameters that describe the mass distributions, as well as a set that describe the spins, were inferred probabilistically using the hierarchical framework outlined in Sec. 6.1.

First, we simulate 2500 population realisations, where for each realisation we randomly choose a single posterior sample from the distributions of the population hyperparameters. For each of the 2500 realisations, we construct the resulting probability distributions for m_1 , q , χ_i , and θ_i , and then draw 2500 samples from each distribution. We note that we draw values for χ_1 and χ_2 separately from the distribution for χ_i in this step, as well as θ_1 separately to θ_2 , so although the GWTC-3 hypermodel assumes primary and secondary spins are drawn from the same underlying population, we obtain slightly different final distributions due to sampling them separately. For each parameter and each population realisation, we then construct a one-dimensional kernel density estimator (KDE) to obtain a smooth posterior predictive probability distribution (PPD). We also use the spin samples to obtain draws of χ_p and χ_{eff} for each realisation, and construct KDEs also for these. We therefore have 2500 population realisations, each with a smooth probability distribution for m_1 , q , χ_1 , χ_2 , θ_1 , θ_2 , χ_p and χ_{eff} . We then construct the interpolated median PPD for each parameter, from which we then draw the binary parameters m_1 , q , χ_1 , χ_2 , θ_1 and θ_2 . The resulting median PPDs (teal) and 90% credible intervals (grey) for the spin parameters are shown in Figs. 6.4 and 6.5.

To obtain the values of the hyperparameters in the **Gaussian Spin** and **Beta Magnitudes**

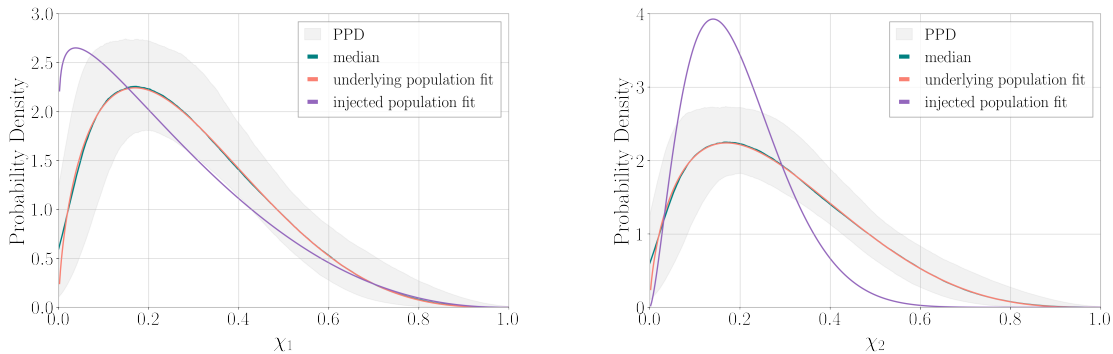


Figure 6.4: Median PPDs (teal), 90% credible intervals (filled grey) for the primary and secondary spin magnitudes distribution χ_1 (left), χ_2 (right), drawn from the inferred population from GWTC-3. We also plot the best-fitting Beta distributions to the median PPDs in orange, and the Beta distribution which best fits our sample population of 20 binaries in purple. We note that due to small number statistics of our injected population, the purple Beta distributions look significantly different to the underlying population distributions, but that the best-fitting Beta distributions in orange match the median PPD very closely.

which describe this underlying population distribution, we perform fits of the median PPDs to the functional forms of these hypermodels. For the spin magnitudes, we use `SciPy curve fit` to fit the median PPD for χ_1 (shown in teal in the left panel of Fig. 6.4) to a Beta distribution (Eq. (6.14)). We do the same for the secondary spin magnitude, and plot the resulting Beta distributions in orange in Fig. 6.4. We note that these Beta distribution fits match the underlying PPD very closely, and so are a good description of the underlying spin population. For the `Gaussian Spin` hypermodel, we perform a 2D fit on the 100,000 initial sample of binaries taken from the underlying population, as 1D fits to the median PPDs shown in teal in Fig. 6.5 will not provide a value for the correlation parameter ρ . We note since these binaries are drawn from the median PPDs plot in teal in Fig. 6.5, the histograms of χ_p and χ_{eff} values closely match these distributions. We fit the 2D histogram in $\chi_p, \chi_{\text{eff}}$ to a 2D Gaussian, and plot the resulting marginalised distributions for χ_p and χ_{eff} in orange in Fig. 6.5. We note that unlike for the `Beta Magnitudes` fit where the orange distribution closely matches the median PPD, here we see significant deviations especially in χ_p , which suggests that our underlying population

of binaries is not well-modelled by a 2D Gaussian. The values for the hyperparameters given as a result of these fits are given in column 3 of Tab. 6.4, and we hereafter will refer to them as the underlying population values.




Figure 6.5: Median PPDs (teal), 90% credible intervals (filled grey) for the effective spin parameter distributions χ_{eff} (left), χ_p (right), drawn from the inferred population from GWTC-3. We also plot the best-fitting 2D Gaussian to the underlying population in orange, and the 2D Gaussian which best fits our sample population of 20 binaries in purple. We note that due to small number statistics of our injected population, the purple distributions look significantly different to the underlying population distributions. Additionally, the underlying population is not well represented by a 2D Gaussian, as shown particularly in the right panel where the best-fitting orange curve for χ_p does not match the underlying median PPD in green.

In addition to the intrinsic binary parameters $\vec{\theta}$ that characterise each binary, BBH mergers are distributed in the expanding universe, which can be probed with GW observations. This also provides a unique opportunity to probe the evolution of the BBH merger \mathcal{R} as a function of redshift. Due to the limited sensitivity range of current GW detectors, the merger rate is commonly modelled as a power law in redshift z ,

$$\mathcal{R}(z) \propto (1+z)^\kappa, \quad (6.35)$$

and we take the value of κ to be 2.9, which is the median value inferred from the GWTC-3 population analysis [35]. We draw binary samples of redshift from this power law distribution between a minimum of $z = 0$ and maximum $z = 5$, where we choose $z = 5$

6. CONSTRAINING PRECESSION IN THE POPULATION

Binary	$m_{1,\text{det}}$ [M_\odot]	$m_{2,\text{det}}$ [M_\odot]	χ_1	χ_2	θ_1 [deg]	θ_2 [deg]	χ_p	χ_{eff}	ϕ_{12} [deg]
1	15.07	13.88	0.02	0.13	90.34	107.23	0.11	-0.02	23.77
2	15.80	14.09	0.54	0.29	52.51	146.39	0.42	0.05	279.72
3	70.97	58.34	0.25	0.16	148.12	136.78	0.13	-0.18	210.27
4	126.60	75.99	0.47	0.22	160.73	50.44	0.15	-0.23	10.96
5	200.11	179.21	0.23	0.07	140.01	148.23	0.15	-0.13	214.60
6	40.69	33.80	0.49	0.26	84.22	125.18	0.49	-0.05	236.13
7	73.24	60.42	0.12	0.23	38.86	83.51	0.18	0.06	355.95
8	117.17	93.77	0.00	0.42	3.56	138.87	0.21	-0.14	281.21
9	27.84	26.44	0.65	0.08	42.50	24.35	0.43	0.28	90.06
10	13.50	9.67	0.60	0.10	95.30	58.70	0.59	-0.01	40.85
11	11.86	10.54	0.06	0.13	131.54	101.73	0.11	-0.04	300.64
12	71.75	64.99	0.18	0.07	66.68	10.61	0.16	0.07	160.57
13	98.99	79.92	0.19	0.23	84.09	52.66	0.19	0.07	227.53
14	56.74	54.58	0.10	0.06	94.17	144.35	0.09	-0.04	79.30
15	27.33	25.12	0.18	0.07	141.56	80.55	0.11	-0.07	2.99
16	22.98	19.08	0.24	0.16	122.39	121.52	0.20	-0.12	310.89
17	73.17	62.93	0.24	0.34	51.26	16.12	0.18	0.23	246.86
18	116.95	92.22	0.15	0.37	34.85	63.43	0.25	0.14	291.55
19	86.52	84.31	0.05	0.09	95.84	155.11	0.05	-0.05	272.45
20	65.18	43.18	0.13	0.33	54.56	71.90	0.20	0.08	120.36

Table 6.2: Injected intrinsic binary parameters, including detector frame masses, and spins, for the 20 binaries in our population. All spin and orientation-related parameters are defined at a reference frequency of $f_{\text{ref}} = 5\text{Hz}$.

to limit the number of samples in our population which would be highly redshifted and therefore out-of-band at a low frequency of 20Hz. We then convert these samples into luminosity distance using (see eg. Ref. [248])

$$d_L(z) = \frac{c(1+z)}{H_0\sqrt{\Omega_k}} \sinh \sqrt{\Omega_k} \int_0^{z'} \frac{dz'}{\sqrt{\Omega_m(1+z')^3 + \Omega_k(1+z')^2 + \Omega_\Lambda}}, \quad (6.36)$$

where c is the speed of light, H_0 is the present day Hubble constant, and Ω_k , Ω_m and Ω_Λ are the dimensionless densities of spacetime curvature, mass and the cosmological constant. We use Planck15 [249] cosmology for these values.

For the remaining binary parameter, we sample uniformly in ϕ_1 (the phase of the primary BH at the reference frequency of 5Hz), $\phi_{12} = \phi_2 - \phi_1$, right ascension α , declination δ , polarization angle ψ and the cosine of the binary inclination $\cos(\iota)$. We also impose a

Binary	α [deg]	δ [deg]	z	d_L [Mpc]	ψ [deg]	θ_{JN} [deg]	ρ_{network}
1	17.09	-36.97	0.29	1543.28	42.05	43.00	146.2
2	100.86	-17.76	0.42	2397.09	131.23	20.08	164.4
3	204.84	-40.39	1.06	7379.77	30.23	47.83	96.2
4	321.44	-41.78	1.37	10063.66	5.61	155.53	107.1
5	26.05	-59.18	2.67	22643.24	153.72	144.82	121.0
6	253.11	26.50	0.27	1468.73	83.79	66.30	238.7
7	67.31	-30.86	1.62	12316.62	149.45	12.37	118.2
8	18.68	-70.56	2.66	22529.88	129.72	8.84	94.3
9	342.16	43.80	0.56	3351.27	86.20	151.48	194.1
10	52.38	-36.47	0.34	1888.33	160.37	123.54	113.4
11	183.26	74.09	0.24	1258.24	93.08	71.96	74.2
12	139.40	64.20	1.06	7335.08	26.36	110.62	120.4
13	272.53	7.66	1.67	12833.73	126.38	168.28	131.8
14	311.87	-37.82	0.74	4733.08	135.29	123.54	113.8
15	3.29	34.61	0.13	672.59	5.31	22.66	720.6
16	73.98	-40.63	0.53	3143.11	163.05	40.74	144.7
17	121.88	64.97	0.95	6457.91	9.96	68.57	125.1
18	179.86	-30.46	2.18	17754.75	170.93	23.66	44.6
19	204.38	-57.86	1.60	12214.68	56.24	144.51	116.7
20	248.23	-3.81	0.89	5927.14	13.18	137.98	130.5

Table 6.3: Injected extrinsic binary parameters, including network SNR ρ_{network} for the 20 binaries in our population. All spin and orientation-related parameters are defined at a reference frequency of $f_{\text{ref}} = 5\text{Hz}$.

minimum/maximum component mass of $5M_{\odot}$ and $87M_{\odot}$ respectively, which are the median values of $m_{\text{min}}, m_{\text{max}}$ as inferred from GWTC-3, and we wish to only sample masses which represent a realistic mass range of the BBH population.

We keep drawing binary parameters $\vec{\theta}_i$ until we reach $i = 100,000$ binaries which satisfy the mass cut conditions, and form a representative sample of the population inferred from GWTC-3 using the GWTC-3 hypermodel. We then compute the SNR for each binary in each detector, and impose a network SNR cut of $\rho_{\text{network}} \geq 12$, which is the chosen detectability threshold. Additionally, we wish to avoid binaries whose network SNR is dominated by CE as this could result in a poor sky localisation and may lead to poorly constrained posteriors due to possible degeneracies between intrinsic and extrinsic parameters for precessing binaries. Therefore, we impose an additional criterion that either $\rho_{\text{L+}} \geq 6$, or $\rho_{\text{V}} \geq 6$, or both. Of the 100,000 initial sample, we find that 524 meet these

mass and SNR criteria.




Figure 6.6: Distribution of the detector frame total mass M in solar masses M_\odot and χ_p for the 20 binaries selected from the population, coloured by the inclination angle θ_{JN} . These binary parameters were drawn from the inferred population from GWTC-3 using the GWTC-3 hypermodel, and are subject to mass and SNR cuts as described in the main text.

Even though more than 500 binaries satisfy our cuts, due to computational reasons we randomly select only 20 of these binaries to form our final subset. The parameters of these binaries are shown in Tabs. 6.2 and 6.3. The distribution of the detector-frame total mass M and χ_p for these 20 selected binaries is shown in Fig. 6.6, coloured by the inclination angle θ_{JN} . We see that our selection of binaries covers a wide mass range between 22 M_\odot and 380 M_\odot , with a preference towards heavier binaries than seen in our current population as future GW detectors will allow us to observe binary mergers at higher redshifts (see Eq. (2.66) for the relation between detector frame masses and redshift). We also note that our binary sample includes a range of inclination angles between almost perfectly aligned to anti-aligned with the orbital angular momentum, and that the χ_p -distribution shows a preference toward moderate precession as expected given the population, with a few binaries displaying more extreme precession of $\chi_p \sim 0.5$.

Fig. 6.7 shows the number of precession cycles calculated from Eq. (6.18) for each of the 20 selected binaries for the three different values of f_{low} . The x -axis is ordered such




Figure 6.7: Number of precession cycles $\alpha(f_{\text{low}}, f_{\text{ISCO}})/2\pi$ for the 20 selected binaries whose parameters are shown in Tabs. 6.2 and 6.3, computed as in Eq. 6.18. We show the number precession cycles for each binary computed from three different values of f_{low} : 20Hz in purple, 10Hz in teal, and 5Hz in orange. As the value of f_{low} decreases, the number of precession cycles observed increases for each binary.

that the numbers of precession cycles increases from right to left, but the binary indexes correspond to those in Tabs. 6.2 and 6.3. As expected, the number of precession cycles increases as f_{low} decreases, and the binaries with smaller total masses spend more time in-band and will therefore generally display more precession cycles.

We caution that due to small number statistics, these 20 binaries will have a slightly different distribution spin distribution than the underlying population from which they were sampled. Therefore, we perform functional fits to the spin parameters of our 20 binaries to determine the hyperparameter values that best describe them. Concretely, and similarly to the fits for the underlying population median PPD, we use the `Scipy curve fit` function to fit Beta distributions to both χ_1 and χ_2 to obtain hyperparameters for the `Beta Magnitudes` hypermodel. For the `Gaussian Spin` hypermodel, we use the same technique as for the underlying population fits of a 2D Gaussian fit to the 2D histogram of χ_p , χ_{eff} . The hyperparameters we obtain are given in the fourth column of Tab. 6.4. We note that for the spin magnitudes, the values for our injected population are similar to those of the underlying population for three out of the four hyperparameters, which

6. CONSTRAINING PRECESSION IN THE POPULATION

are shown in the third column of the same table. Additionally, we also fit our selected 20 binaries to a bivariate Gaussian in χ_p and χ_{eff} , and obtain the corresponding hyperparameters for the injected population also given in Tab. 6.4. These values vary slightly from those of the underlying population. In particular, we find that the hypermodel means μ_{χ_p} and $\mu_{\chi_{\text{eff}}}$ are shifted relative to the population they were initially drawn from, with our 20 binary population showing on average less precession, and more correlation between χ_p and χ_{eff} . The effect of the differences in these hyperparameter values can be seen in Figs. 6.4, 6.5 where we plot the injected population distributions in purple alongside the best-fit underlying population distributions in orange. Notably, the peak of the primary spin magnitude is shifted more towards lower values, and the width of the secondary spin distribution is narrower than for the underlying population. The distribution on χ_p is also narrower, and shifted towards lower precession values, and the distribution of χ_{eff} becomes slightly broader and more negative.

Hyper-model	Hyper-parameter	Population value	Injected value	Prior	Recovered value		
					5Hz	10Hz	20Hz
Beta Magnitudes	α_{χ_1}	1.63	1.09	U(1, 7)	$1.86^{+0.72}_{-1.24}$	$2.22^{+1.31}_{-0.94}$	$3.87^{+1.77}_{-1.72}$
	α_{χ_2}	1.63	2.54	U(1, 7)	$3.52^{+1.85}_{-1.67}$	$4.40^{+2.01}_{-2.15}$	$4.94^{+1.74}_{-2.27}$
	β_{χ_1}	4.07	3.29	U(1, 10)	$4.72^{+3.15}_{-1.99}$	$6.06^{+3.18}_{-2.74}$	$7.64^{+2.11}_{-3.43}$
	β_{χ_2}	4.06	10.42	U(1, 25)	$17.60^{+6.63}_{-8.62}$	$16.81^{+6.97}_{-8.33}$	$18.91^{+5.32}_{-8.97}$
Gaussian	$\mu_{\chi_{\text{eff}}}$	0.02	-0.06	U(-1, 1)	$-0.01^{+0.07}_{-0.06}$	$-0.01^{+0.07}_{-0.06}$	$0.00^{+0.05}_{-0.05}$
	$\sigma_{\chi_{\text{eff}}}$	0.13	0.19	U(0, 1)	$0.16^{+0.05}_{-0.08}$	$0.16^{+0.05}_{-0.08}$	$0.14^{+0.04}_{-0.07}$
	μ_{χ_p}	0.24	0.15	U(0, 1)	$0.18^{+0.15}_{-0.10}$	$0.20^{+0.16}_{-0.10}$	$0.29^{+0.10}_{-0.06}$
	σ_{χ_p}	0.15	0.06	U(0, 1)	$0.22^{+0.09}_{-0.16}$	$0.23^{+0.10}_{-0.17}$	$0.12^{+0.05}_{-0.14}$
	ρ	0.00	0.44	U(-0.75, 0.75)	$0.33^{+0.52}_{-0.35}$	$0.35^{+0.52}_{-0.34}$	$0.05^{+0.47}_{-0.46}$

Table 6.4: Injected and recovered values of the population hyperparameters for each of the **Beta Magnitudes** and **Gaussian spin** models, for each of the f_{low} cutoffs of 5Hz, 10Hz and 20Hz. The ‘population values’ for the hypermodels are the result of a fit to the population posterior predictive probability distribution from which we drew our population of 20 binaries. The injected hyperparameter values are the result of a functional fits of the corresponding hypermodel to the 20 binaries which were analysed. The recovered values show the median for each hyperparameter, as well as the 90% credible interval.

6.3.2 Selected Individual Binary Parameter Estimation Results

In this subsection, we present results of the parameter estimation for two of the binaries. Figs. 6.8 and 6.9 show the mass and spin parameter posteriors respectively for binary 6 in Tabs. 6.2 and 6.3. This binary lies towards the lower end of the total mass range covered by our 20 binaries, has a relatively large value of χ_p at 0.49, and the number of precession cycles that this binary undergoes lies towards the middle of our sample. We show the 90% 2D credible intervals and 1D posterior distributions for $f_{\text{low}} = 5\text{Hz}$ in orange, 10Hz in teal and 20Hz in purple, with the true injected values from Tabs. 6.2 and 6.3 marked in black. We note that all three of the analyses recover the mass parameters very well, with the peaks of the 1D posterior distributions centred on the true injected values. The most noticeable difference between the three distributions lies in the width of the 90% credible intervals, particularly in the chirp mass \mathcal{M} , and the total mass M . The 1D marginalised posterior distribution for \mathcal{M} has a 90% credible interval of $0.11 M_{\odot}$ for $f_{\text{low}} = 20\text{Hz}$, which shrinks dramatically to $0.02 M_{\odot}$ at 10Hz, and then $0.01 M_{\odot}$ at 5Hz, leading to an accuracy improvement of an order of magnitude between 20Hz and 5Hz. Similarly, the total mass M has a width of $0.27 M_{\odot}$ for $f_{\text{low}} = 20\text{Hz}$, which then reduces by 50% to $0.13 M_{\odot}$ as f_{low} decreases to 10Hz, and reduces further to $0.12 M_{\odot}$ at 5Hz, a total reduction of 55% in uncertainty. The component masses m_1 , m_2 and mass ratio q show more modest improvements, with the 90% credible width reducing between 20Hz and 5Hz by 22% for m_1 and 30% for m_2 , and q reducing by 23% between 20Hz and 5Hz.

The posterior distributions for the spin parameters for this binary are shown in Fig. 6.9, similarly showing 20Hz in purple, 10Hz in teal, 5Hz in orange, and the injected values in black. Again, we observe a clear overall improvement in the 90% credible interval widths of the recovered parameters as f_{low} decreases. The 90% credible width of the marginalised posteriors for the spin magnitudes χ_1 and χ_2 are 0.184 and 0.157 at $f_{\text{low}} = 20\text{Hz}$ respectively, which then reduces by 58% for χ_1 and 51% for χ_2 as f_{low} is lowered to 10Hz, and reduces to 34% (χ_1) and 40% (χ_2) of the 20Hz width at 5Hz. The spin tilts also reduce from a width of 9.3 and 33.3 degrees for θ_1 and θ_2 at 20Hz, to just 6 and 12

degrees at 5Hz. The largest improvements can be seen in the effective spin parameters χ_p and χ_p whose widths reduce by 76% and 67% respectively between 20Hz and 5Hz.

Furthermore, we find that the median values of the spin parameters are recovered more accurately as f_{low} is decreased. For example, the median posterior value for χ_p at $f_{\text{low}} = 20\text{Hz}$ is 0.46, which is an offset (bias) of 0.03 compared to the injected value of 0.49 in Tab. 6.2. This reduces to 0.02 at 10Hz, where the recovered median value is 0.51, and further to 0.01 at 5Hz, where the median is 0.50. The median offset in χ_1 also decrease, from 0.04, to 0.02, to 0.01, and we see similar trends across all the spin parameters. Our results verify that the combination of higher SNR and more observed precession cycles leads to significant improvements in the recovered spin posteriors.

In Figs. 6.10 and 6.11 we show results for a second binary, binary 5 in Tabs. 6.2 and 6.3, which is the heaviest of all binaries in our sample and shows the fewest number of precession cycles in Fig. 6.7. While we find a relatively good recovery of the injected mass values at all values of f_{low} as seen in Fig. 6.10, there are improvements especially in the chirp mass \mathcal{M} . The component masses m_1 and m_2 show modest improvements, with the 90% credible interval width of m_1 and m_2 decreasing from $11.8 M_\odot$ and $8.8 M_\odot$ at 20Hz respectively, to $7.2 M_\odot$ and $7.5 M_\odot$ at 10Hz, an uncertainty reduction of 39% in m_1 and 15% in m_2 . At 5Hz, the m_2 reduces slightly further to $7.2 M_\odot$, while the m_1 stays the same as at 10Hz. We see similar trends in the other mass parameters, namely that there is a significant reduction in the posterior widths between 20Hz and 10Hz, and then little change between 10Hz and 5Hz. The largest difference between 10Hz and 5Hz can be seen in the chirp mass \mathcal{M} , where the 20Hz 90% credible interval width of $4.1 M_\odot$ reduces to $2.6 M_\odot$ at 10Hz, a decrease of 40%. Similarly, the total mass width decreases by 37% from $9.7 M_\odot$ at 20Hz to $5.8 M_\odot$ at 10Hz. Lastly, the mass ratio q width reduces by 33% from 0.09 at 20Hz to 0.06 at 10Hz.

Lastly, we note that similarly to binary 6, the biases in the recovered median spin parameters decreases as f_{low} decreases. A clear example is shown in Fig. 6.11 for the primary spin tilt θ_1 . The difference between the true injected value and the median




Figure 6.8: 90% 2D credible intervals and 1D posterior distributions for the detector-frame mass parameters of representative binary 6 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$ in purple, 10Hz in teal, and 5Hz in orange, and plot the true injected values as shown in Tab. 6.2 in black. As f_{low} decreases, the injected values are better recovered with lower uncertainty.

posterior value decreases from 17 degrees at 20Hz, to 6 degrees at 10Hz, down to just 2 degrees at 5Hz. Similarly, the bias in χ_p reduces from 0.2 at 20Hz, to 0.05 at 10Hz, and 0.04 at 5Hz. We also see a decreasing bias in the secondary spin magnitude, from 0.26 at 20Hz, down to 0.1 at 10Hz, and 0.05 at 5Hz. An exception to this trend is the secondary tilt angle θ_2 , for which none of the three f_{low} values accurately recovers the injected value at the peak of the 1D marginalised posterior distributions.

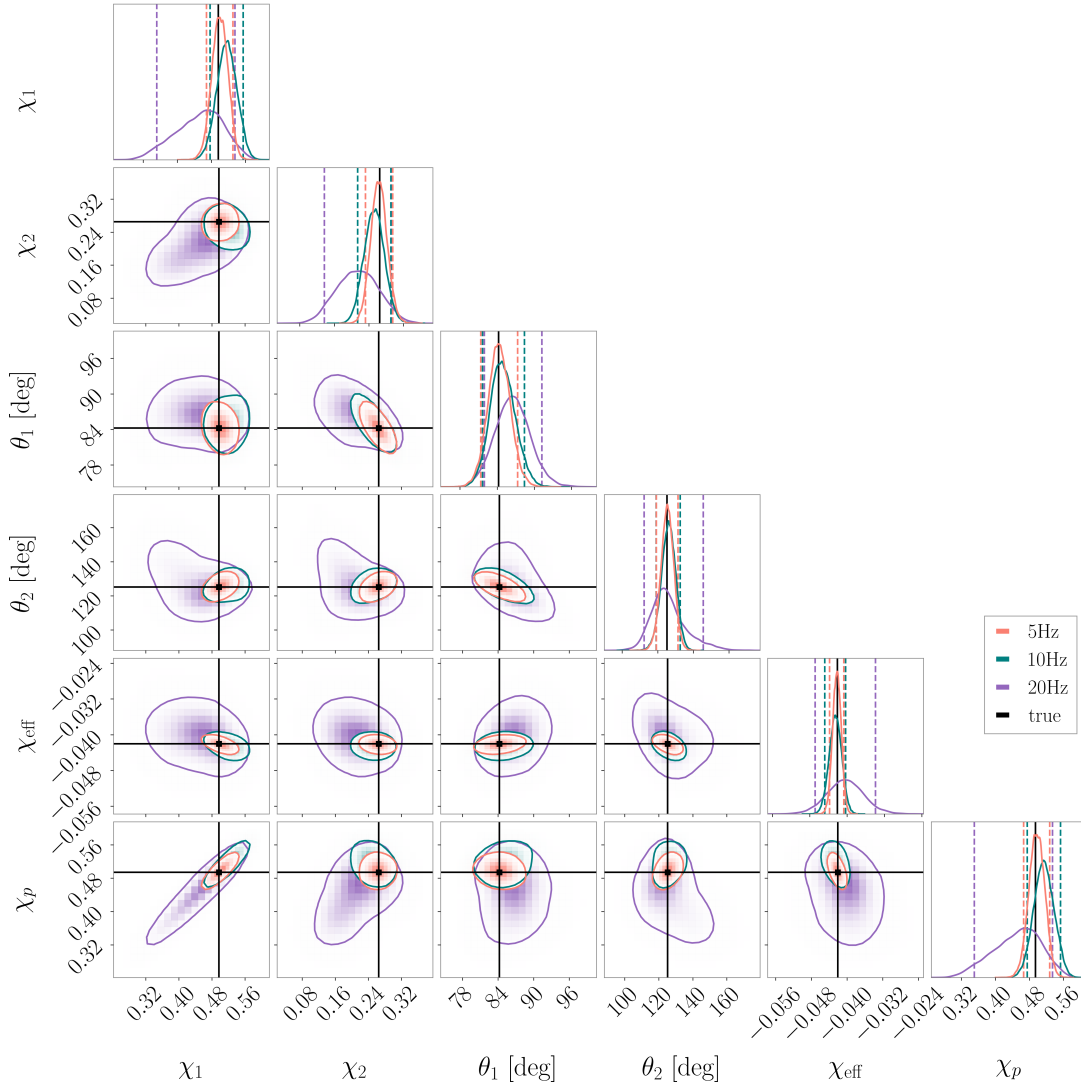


Figure 6.9: Parameter estimation results for spin parameters for representative binary 6 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$ in purple, 10Hz in teal, and 5Hz in orange, and plot the true injected values as shown in Tab. 6.2 in black. The injected and recovered spin values for the are specified at a reference frequency of 5Hz. As f_{low} decreases, the injected values are better recovered with lower uncertainty.

In Fig. 6.12 we show the bias and 90% credible interval widths for χ_{eff} in the left panel, and χ_p on the right for all 20 binaries in our sample. In order to show all 20 binaries on the same figure, we subtract from each posterior sample the true injected value of the parameter for the binary in question. As an example, for binary 3 which has a true injected χ_p value of 0.13 as shown in Tab. 6.2, we subtract 0.13 from all the χ_p posterior samples for this binary. We call this difference $\Delta\chi_p$, which is plotted in the right panel

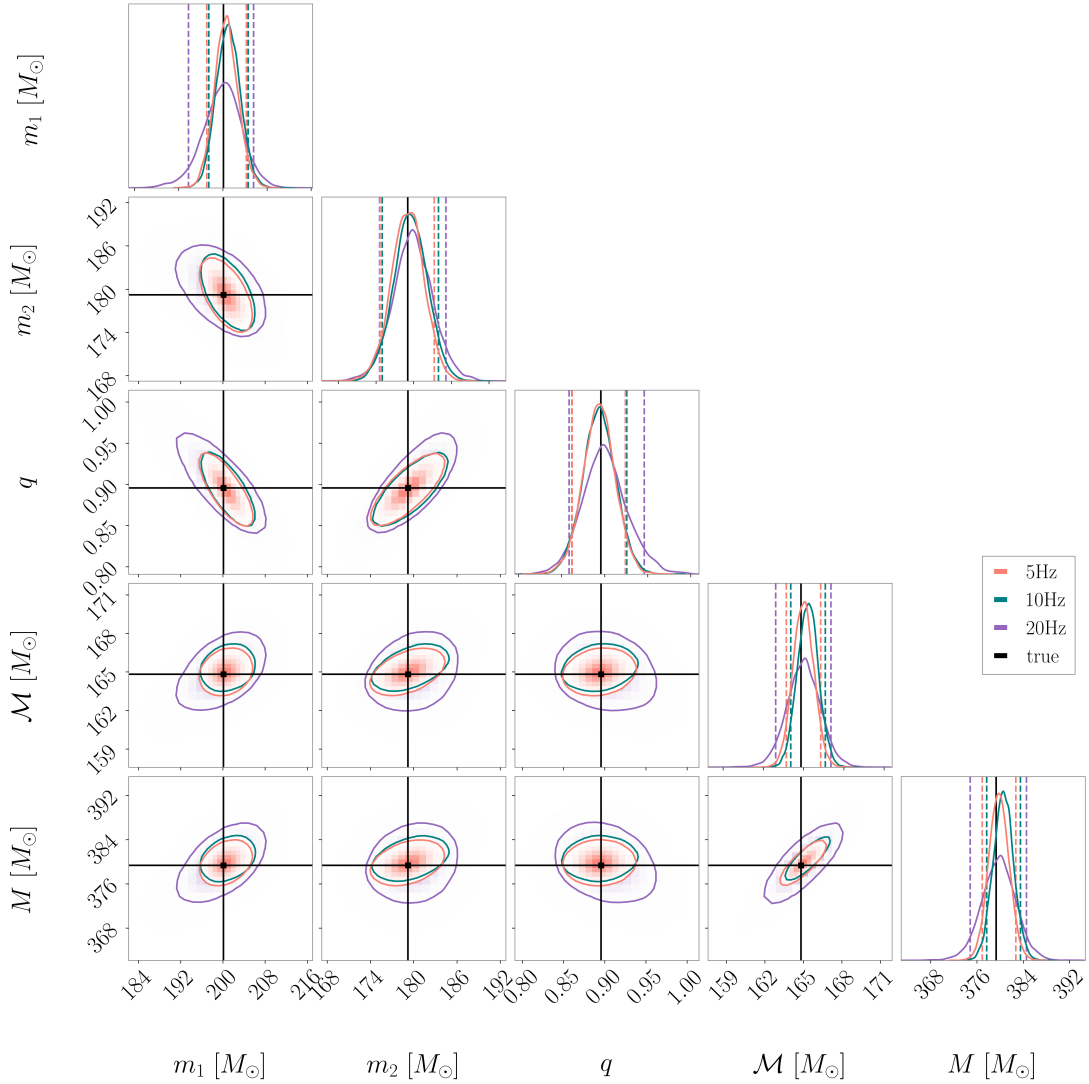


Figure 6.10: 90% 2D credible intervals and 1D posterior distributions for the detector-frame mass parameters of binary 5 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$ in purple, 10Hz in teal, and 5Hz in orange, and plot the true injected values as shown in Tab. 6.2 in black. As f_{low} decreases, the injected values are better recovered with lower uncertainty.

of Fig. 6.12. Analogously, we define $\Delta\chi_{\text{eff}}$ which is plotted on the left. We plot the 90% credible intervals on $\Delta\chi_{\text{eff}}$ and $\Delta\chi_{\text{eff}}$ as shaded bands. The coloured points show the difference between the maximum a posteriori probability samples (as defined in Sec. 2.5) and the true injected values.

As in previous figures, we show the analyses with $f_{\text{low}} = 20\text{Hz}$ in purple, 10Hz in teal and 5Hz in orange, and analogously to Fig. 6.7 we have ordered the binaries such that the




Figure 6.11: Parameter estimation results for spin parameters for binary 5 for each of the 3 f_{low} values. We show results for $f_{\text{low}} = 20\text{Hz}$ in purple, 10Hz in teal, and 5Hz in orange, and plot the true injected values as shown in Tab. 6.2 in black. The injected and recovered spin values for the are specified at a reference frequency of 5Hz . As f_{low} decreases, the injected values are better recovered with lower uncertainty.

90% credible widths for the 5Hz analyses increase from left to right for ease of reading.

We first note that as f_{low} is lowered, we see a general trend of the 90% credible interval widths decreasing for both the χ_p and χ_{eff} posteriors across the 20 binaries. This effect is more pronounced in χ_{eff} than χ_p , particularly for binaries where χ_{eff} is negative such as binaries 3, 4, 5 and 8. These binaries also each have relatively high total masses. This suggests that for binaries with shorter GW signals in band due to their high mass and

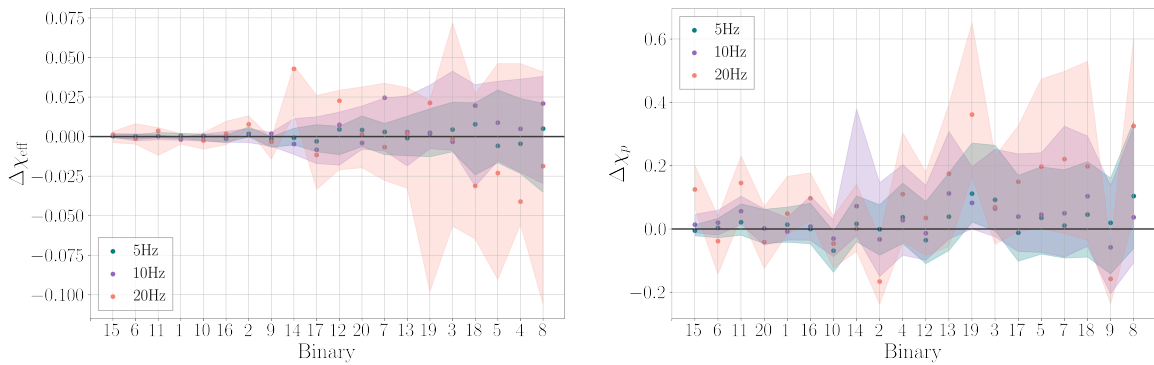


Figure 6.12: Left: $\Delta\chi_{\text{eff}}$ for the 20 selected binaries, defined as the deviation of the posterior samples from the true injected χ_{eff} value for each binary. Right: Results of $\Delta\chi_p$ for the individual 20 binaries, defined as the deviation of the posterior samples from the true injected χ_p value for each binary. We show the 90% credible interval widths as shaded bands, and the maximum posterior points as coloured markers, for $f_{\text{low}} = 20\text{Hz}$ in orange, 10Hz in purple, and 5Hz in green. We see a general trend that as f_{low} decreases, the widths of the 90% credible interval also decrease.

negative χ_{eff} , having a lower cutoff frequency f_{low} is even more crucial to precisely measure the binary spins. For χ_p in the right panel there is not as clear a trend between the value of χ_p and the recovered width of the posteriors. But we also note that for all binaries except binary 19, which has the smallest χ_p value of our population (see Tab. 6.2), the true injected value lies within the 90% credible interval for all analyses with $f_{\text{low}} = 5\text{Hz}$. The 10Hz analyses has a single binary where the true value does not lie in the 90% credible interval, binary 11, which again has a relatively low χ_p value at 0.11 as seen in Tab. 6.2. However, as f_{low} increases to 20Hz, this number increases to 6 binaries where the true value lies outside the 90% credible interval. These 6 binaries include binaries 19, 11, and 15, all of which have relatively low χ_p values. This suggests that especially for binaries where there is little precession, a lower minimum frequency of f_{low} allows us to measure χ_p precisely and without bias.

We also show in Fig. 6.13 $\Delta\chi_1$ (left) and $\Delta\chi_2$ for the 20 binaries, defined analogously as for $\Delta\chi_p$ and $\Delta\chi_{\text{eff}}$. Again we show results for 20Hz in orange, 10Hz in purple and 5Hz in green, and plot the 90% credible intervals as shaded bands and the maximum posterior probability points as coloured markers. We order the binaries such that $\Delta\chi_1$ and $\Delta\chi_2$ for

the 5Hz maximum posterior points increases from left to right. While the trend is not as clear for the spin magnitudes as for the effective spin parameters shown in Fig. 6.12, we do see that in general, the values of $\Delta\chi_p$ and $\Delta\chi_{\text{eff}}$ at the maximum posterior points are smallest for $f_{\text{low}} = 5\text{Hz}$, and that the 90% credible interval widths decrease with lower minimum frequency cutoff.

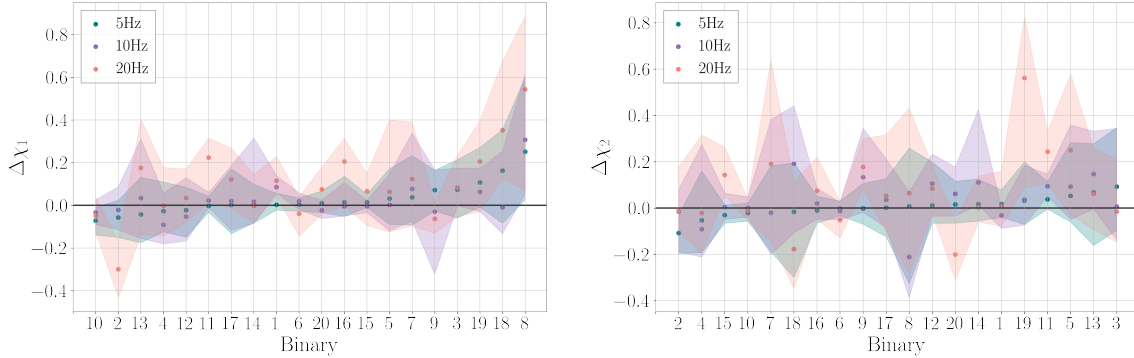


Figure 6.13: Left: $\Delta\chi_1$ for the 20 selected binaries, defined as the deviation of the posterior samples from the true injected χ_1 value for each binary. Right: Results of χ_2 for the individual 20 binaries, defined as the deviation of the posterior samples from the true injected χ_2 value for each binary. We show the 90% credible interval widths as shaded bands, and the maximum posterior points as coloured markers, for $f_{\text{low}} = 20\text{Hz}$ in orange, 10Hz in purple, and 5Hz in green. We see a general trend that as f_{low} decreases, the widths of the 90% credible interval also decrease.

In summary, in this section we have shown that for our sample of 20 binaries, by decreasing the value of f_{low} we gain tighter constraints on our recovered spin parameters for each individual binary. We have shown that the 90% credible interval widths for the spin parameters decrease by reducing the minimum frequency cutoff, and that biases in the recovered spin parameters also decrease. In the next section we determine if these tighter constraints and better recovered spins for the individual binaries will result in more accurate inference of the spin distribution across the entire population.

6.3.3 Population Inference Results

In the previous subsection we presented parameter estimation results for individual binaries, showing that as f_{low} decreases, the accuracy with which we can measure the mass

and spin parameters improves. Now we assess whether this improvement can also be seen when estimating the hyperparameters that describe underlying population of binaries. We make use of the hierarchical Bayesian formalism outlined in Sec. 6.1, and choose two different hypermodels for our population spin distributions: i) the **Beta Magnitudes** hypermodel as described in Subsec. 6.1.1, and ii) the **Gaussian** hypermodel of Subsec. 6.1.2. We then perform three analyses for each hypermodel, corresponding to each of the three f_{low} values from which we performed our individual binary estimation as described in Sec. 6.3.2. We use the `gwpopulation` [250] Python package to perform the hierarchical analysis, for the sampling we use `Dynesty` within `Bilby` similarly to the individual event analyses in Sec. 6.2 and Subsec. 6.3.3. The priors for the population hyperparameters are detailed in Tab. 6.4, and we use (rounded to the nearest integer) as our prior bounds for the **Beta Magnitudes** hypermodel the minimum and maximum samples for the hyperparameters contained within the **GWTC-3** data release [251] except for β_{χ_2} , where we raise our prior upper bound to 25. This accounts for the fact that our injected population of 20 binaries has a β_{χ_2} value of 10.42, and so an upper prior boundary of 10 would not be sufficient. For the **Gaussian** hypermodel we use the same prior bounds as used in Ref. [35]. In the remainder of the section, we present our results for the population inference, first the **Gaussian Spin** hypermodel results in Subsubsec. 6.3.3 and then the **Beta Magnitudes** results in Subsubsec. 6.3.3.

Effective Spin Parameter Results

The results for the **Gaussian** hypermodel for the effective spin parameters are shown in Fig. 6.14, where we show the 1D marginalised posterior distributions for each of the five hyperparameters $\vec{\Lambda} = \{\mu_{\chi_{\text{eff}}}, \sigma_{\chi_{\text{eff}}}, \mu_{\chi_p}, \sigma_{\chi_p}, \rho\}$, as well as the 2D 90% credible contours. As before, we show the results of the analysis with $f_{\text{low}} = 20\text{Hz}$ in purple, 10Hz in teal, and 5Hz in orange. The grey lines mark the true median values of the hyperparameters as taken from the underlying distribution from which we drew our population of 20 binaries, ie. the hyperparameter values shown in column 3 of Tab. 6.4. The black lines mark




Figure 6.14: Results of population inference with the **Gaussian Spin** hypermodel. We show the posteriors on the hyperparameters $\{\mu_{\chi_{\text{eff}}}, \sigma_{\chi_{\text{eff}}}, \mu_{\chi_p}, \sigma_{\chi_p}, \rho\}$ for each of the three f_{low} values, 5Hz in orange, 10Hz in teal and 20Hz in purple. We plot the true injected values of these hyperparameters from our population of 20 binaries in black, and the values taken from the underlying population from which our 20 binaries were drawn in grey.

the median hyperparameter values for the population of 20 binaries, i.e. the injected hyperparameter values shown in column 4 of Tab. 6.4. In Tab. 6.4 we also list the medians and 90% credible interval bounds for the recovered posteriors on the five hyperparameters.

We first note that hyperparameters inferred from the the 20Hz analysis seems to agree better with the true population hyperparameters. However, as discussed below, this is largely driven by the priors. Both the 10Hz and 5Hz analyses are more informative

with the inferred hyperparameters agreeing with the injected population. Whilst this is seemingly biased from the underlying population, as we incorporate more binaries, we would expect both these analyses to converge on the true population hyperparameters quicker. The reason for the 20Hz analysis better recovering the underlying population hyperparameters can be interpreted as follows. With an f_{low} of 20Hz in comparison to 10Hz or 5Hz, we have already seen in Secs. 6.2 and 6.3 that the individual event posteriors on the spin parameters are broader due to a combination of lower SNR and fewer observed precession cycles. Therefore, the results for the individual binaries are prior-driven, and given the priors we have used which are uniform in spin magnitudes and isotropic in direction, the prior on χ_{eff} resembles a Gaussian centred on zero (see eg. Ref. [252]). The prior on χ_p has a peak at around 0.33, and the two effective spin distributions will be uncorrelated. Therefore the prior-driven values for the hyperparameters correspond to $\mu_{\chi_{\text{eff}}} = 0$, $\mu_{\chi_p} = 0.33$, and $\rho = 0$, which lie very close to the recovered median posterior values for the 20Hz results in Tab. 6.4. Therefore, instead of recovering the underlying population hyperparameter values, the 20Hz analysis instead uncovers the expected results from individual binaries where the event posteriors are prior-driven.

We then note that due to being prior-driven, the 20Hz analysis better recovers the hyperparameters for the underlying population from which we drew our 20 binaries, which are shown by the grey lines in Fig. 6.14. For example, the true population ρ value of 0 lies closer to the median posterior value from the 20Hz analysis, which is 0.05, and is comfortably contained within the 1D marginalised posterior 90% credible interval for 20Hz. In contrast, the median posterior values for the 5Hz and 10Hz are 0.33 and 0.35 respectively, and the population value lies just within the 90% credible interval for the 5Hz results, and just outside the interval for the 10Hz. Another example can be seen in the hyperparameter $\mu_{\chi_{\text{eff}}}$, where the median posterior value for 20Hz is 0.0, which better matches the population value of 0.02 than the median recovered values for 5Hz and 10Hz, which both equal -0.01 . We also note that the 90% credible interval for this hyperparameter is narrower for the 20Hz analysis than either the 5Hz or 10Hz.

Therefore the 20Hz analysis appears to recover the underlying population width of the Gaussian on χ_{eff} more accurately and with less uncertainty than the 5Hz or 10Hz analyses. Additionally, while the underlying population values for $\sigma_{\chi_{\text{eff}}}$ and σ_{χ_p} are contained within the 90% credible widths for all three f_{low} values, for both hyperparameters the median posterior value for the 20Hz analysis lies closer to the population value than either the 10Hz or 5Hz.

We now consider how well the analyses are able to recover the hyperparameter values for the population of 20 binaries that we actually injected, which are shown in column 4 of Tab. 6.4. These injected values are shown in black in Fig. 6.14. We find that the 5Hz most accurately recovers the injected value of $\mu_{\chi_p} = 0.15$, as this analysis has a median posterior value of $0.18_{-0.10}^{+0.15}$, compared to the 10Hz median of $0.2_{-0.10}^{+0.16}$, and the 20Hz of $0.29_{-0.06}^{+0.10}$ (though this 20Hz median value is likely prior-driven). Additionally, we find that as f_{low} decreases to 10Hz and 5Hz we more accurately recover the correlation parameter ρ , as the median shifts away from the prior-driven 20Hz value of 0.05 to 0.35 at 10Hz and 0.33 at 5Hz. We note that interestingly, our 90% credible widths upon the recovered hyperparameters stay consistent across the three low frequency values, and that the dominant effect of lowering f_{low} is to move the peak of the posterior distribution towards the injected values, as opposed to reducing the uncertainty on these recovered values with narrower posteriors.

We remark that although the analyses with $f_{\text{low}} = 5\text{Hz}$ and 10Hz do recover the injected population hyperparameters more faithfully, we caveat that the underlying data is not well represented by a bivariate Gaussian. We have already seen in Fig. 6.5 that the underlying population distributions differ significantly from the best-fitting bivariate Gaussian. This therefore places a limitation on the accuracy we can achieve as we would not expect to recover the population distribution perfectly even with significantly increased information.

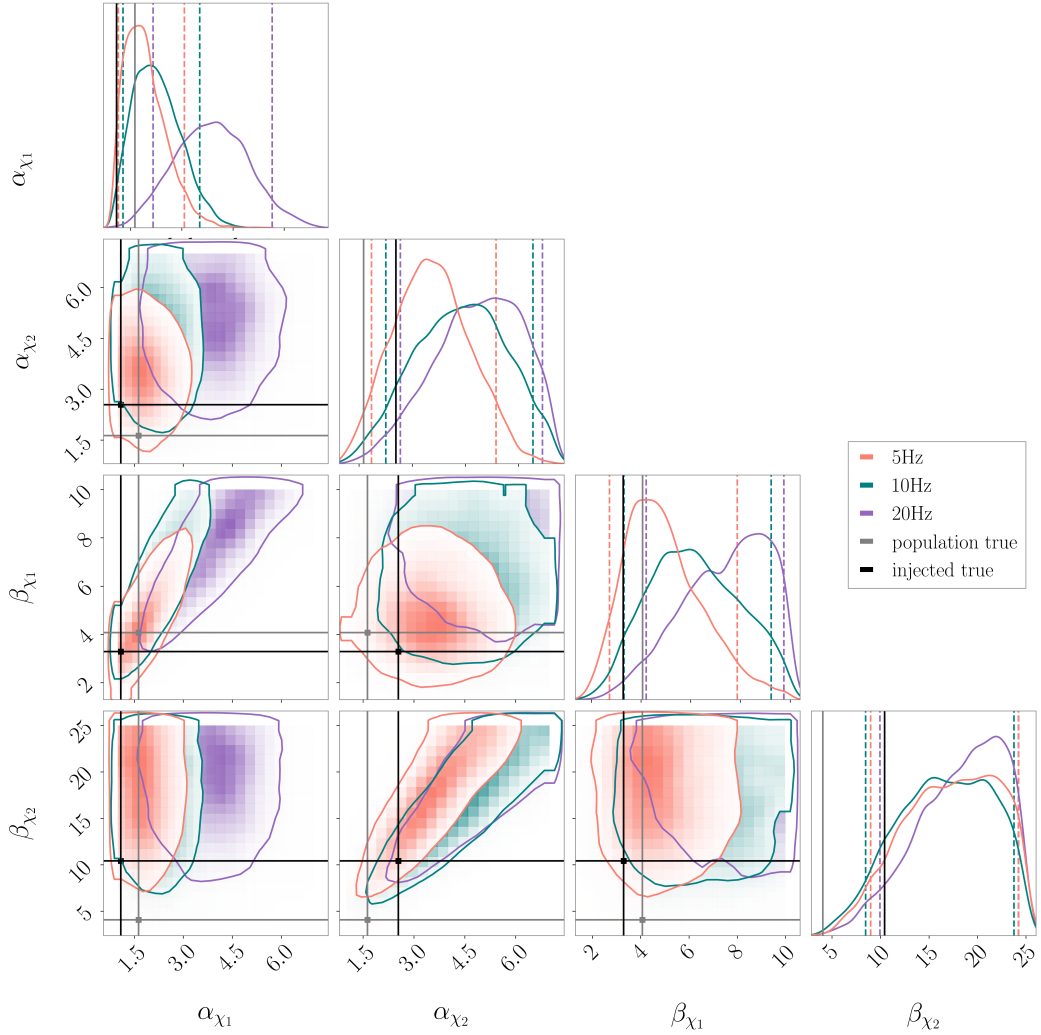


Figure 6.15: Results of population inference with the **Beta Magnitudes** hypermodel. We show the posteriors on the hyperparameters $\{\alpha_{\chi_1}, \alpha_{\chi_2}, \beta_{\chi_1}, \beta_{\chi_2}\}$ for each of the three f_{low} values, 5Hz in orange, 10Hz in teal and 20Hz in purple. We plot the true injected values of these hyperparameters from our population of 20 binaries in black, and the values taken from the underlying population from which our 20 binaries were drawn in grey.

Spin Magnitude Results

The results for the **Beta Magnitudes** hypermodel are summarised in Fig. 6.15, where we show the 1D marginalised posterior distributions and 2D 90% credible interval contours for each of the four hyperparameters $\vec{\Lambda} = \{\alpha_{\chi_1}, \alpha_{\chi_2}, \beta_{\chi_1}, \beta_{\chi_2}\}$, with the 20Hz results in purple, 10Hz in teal and 5Hz in orange. We show the true values from the underlying

population distribution as detailed in column 3 of Tab. 6.4 in grey, and the true injected values of the hyperparameters from our population of 20 binaries in black. The recovered medians and 90% credible interval bounds for each of the three f_{low} values can be found in Tab. 6.4.

We first consider the recovered hyperparameters for the distribution of the primary spin magnitude, α_{χ_1} and β_{χ_1} . For both of these hyperparameters we find that in all our analyses the recovered hyperposterior medians overestimate the true values. The peaks of the 1D marginalised posterior distributions for α_{χ_1} and β_{χ_1} lie consistently to the right of the true injected values shown by the black vertical lines. However, we also see that as f_{low} decreases, this effect reduces. The difference in α_{χ_1} between the median recovered posterior value and the true injected value of 1.09 reduces from 2.78 at $f_{\text{low}} = 20\text{Hz}$, to 1.13 at 10Hz, to 0.77 at 5Hz, a total reduction of 72% in the bias between the 20Hz and 5Hz analyses. Similarly, the bias in β_{χ_1} between the median posterior and true injected value of 3.29 decreases from 4.35 at 20Hz, to 2.77 at 10Hz, to 1.43 at 5Hz, a total reduction of 57% between 20Hz and 5Hz. We also find that the widths of the 90% credible intervals on α_{χ_1} decrease with lower f_{low} . For the 20Hz analysis, the 90% credible interval width is 3.49, which then reduces by 36% to 2.25 as f_{low} lowers to 10Hz. The width at 5Hz is 1.96, which is a total reduction of 44% compared to 20Hz. For β_{χ_1} , the widths of the 90% interval do not show the same decreasing trend, but stay broadly consistent across the different values of f_{low} . The width increases slightly from 5.54 at 20Hz to 5.92 at 10Hz, before reducing slightly to 5.14 at 5Hz. However, we note that the true injected β_{χ_1} of 3.29 is not contained within the 90% credible intervals at 10Hz or 20Hz, but is contained within the interval at 5Hz. Similarly, the true α_{χ_1} injected value of 1.09 is not contained within the 90% credible intervals of either the 20Hz or 10Hz analyses, but is within the 5Hz interval. Therefore, only with an f_{low} of 5Hz do we actually recover the injected primary spin magnitude distribution at the 90% level of the 1D marginalised posterior distributions. If we consider the 2D 90% contour of α_{χ_1} and β_{χ_1} , the true injected value in black is contained within the 90% contours of both

the 5Hz and 10Hz analyses, but not the 20Hz. Therefore we see that a low frequency cutoff of 20Hz in this case is not sufficient to accurately recover the injected population distribution for χ_1 .

We now consider the posteriors for the hyperparameters of the secondary spin magnitude distribution, α_{χ_2} and β_{χ_2} . It can be seen from the definitions of χ_{eff} in Eq. (2.93) and χ_p in Eq. (3.15) that the contributions of the primary and secondary spins upon the waveform are weighted by the BH masses, for example in Eq. (2.93) the secondary spin contribution appears as $m_2\chi_{2\parallel}/M$. Therefore, as $m_2 \leq m_1$, we expect the imprint of the secondary spin upon the waveform to be subdominant compared to the primary spin, and so the secondary spin will be more difficult to constrain for an individual event and on the population level. With this in mind, it is unsurprising that we see wider posteriors on α_{χ_2} and β_{χ_2} than α_{χ_1} and β_{χ_1} across the three values of f_{low} , and less significant improvements as f_{low} decreases. While we do not see as noticeable an improvement in the constraints on these hyperparameters as for the primary spin magnitude hyperparameters, we note that similarly to α_{χ_1} and β_{χ_1} , the tendency of our analysis to overestimate the true value of α_{χ_2} decreases as f_{low} is lowered. The difference between the median recovered posterior value for α_{χ_2} and the true injected value of 2.54 reduces from 2.4 at 20Hz, to 1.86 at 10Hz, to 0.98 at 5Hz, which is a reduction of 54% between 20Hz and 10Hz. We see a slight narrowing of the 1D marginalised posterior width on α_{χ_2} of 12% between 20Hz and 10Hz, though there is a slight increase in width between 20Hz and 10Hz by 4%. We also note that the true injected value of 2.54 is not contained within the 90% credible interval of the 1D marginalised posterior for α_{χ_2} at 20Hz, but it is contained within the 90% interval for the 10Hz and 5Hz analyses. The only hyperparameter in the **Beta Magnitudes** analysis which does not show any noticeable improvement with lower f_{low} is β_{χ_2} . The 1D marginalised posteriors for this parameter remain mostly unconstrained for all values of f_{low} , with wide and comparably-sized 90% credible interval widths across all 3 frequency cutoffs.

Interestingly, we see evidence of a degeneracy between α_{χ_1} and β_{χ_1} , and also between

α_{χ_2} and β_{χ_2} . This is evidenced by the large diagonal 2D contours for these pairs of parameters. We interpret this as follows: for a Beta distribution described by the shape parameters $\alpha \geq 1$ and $\beta \geq 1$, the mean of the distribution is given by $\alpha/(\alpha + \beta)$, and the variance is $\alpha\beta/[(\alpha + \beta + 1)^2(\alpha + \beta + 1)]$. Therefore if the variance of our recovered spin distribution is uncertain, but the mean is well constrained, this places a constraint on the ratio β/α , but not on the individual shape parameters, leading to a degeneracy such as the ones seen in Fig. 6.15. We therefore hypothesise from these results that across the three values of f_{low} , while we are constraining the mean of the spin distribution with this binary population, we are not effectively constraining its width, potentially due to the low number of binaries in our sample.

As mentioned above, and in contrast to the **Gaussian Spin** hypermodel results, as f_{low} decreases the α_{χ_1} and β_{χ_1} hyperparameters for the χ_1 distribution appear to recover the values from underlying population (plotted in grey in Fig. 6.15) as opposed to the true hyperparameter values from the subset of 20 binaries we injected. To investigate this, in Fig. 6.16 we show the median recovered Beta distributions for the primary spin magnitude χ_1 on the left side, and for χ_2 on the right, plotted in orange in the top row for the results where $f_{\text{low}} = 5\text{Hz}$, in teal in the middle row for 10Hz, and in purple on the bottom row for 20Hz. These median Beta distributions are obtained by taking the median posterior samples for α_{χ_i} and β_{χ_i} for each value of f_{low} and evaluating the resulting distributions given by Eq. (6.14). We also show in Fig. 6.16 the Beta distributions for χ_1 and χ_2 given by the hyperparameter values from the injected binaries in black (where the values of α_{χ_i} and β_{χ_i} are given in column 4 of Tab. 6.4), and the values taken from the underlying population from which we chose our 20 binaries in grey (column 3 in Tab. 6.1). Finally, we plot the histograms of the individual event posteriors for χ_1 and χ_2 for each of the 20 binaries, for each value of f_{low} . At $f_{\text{low}} = 20\text{Hz}$, the individual binary posteriors on χ_1 mainly appear to be broad, and as f_{low} is decreased to 10Hz and then 5Hz, these distributions narrow as discussed in Subsec. 6.3.2. However, at 5Hz many of the binaries with moderate values of χ_1 still appear to have broad posterior distributions, in contrast to

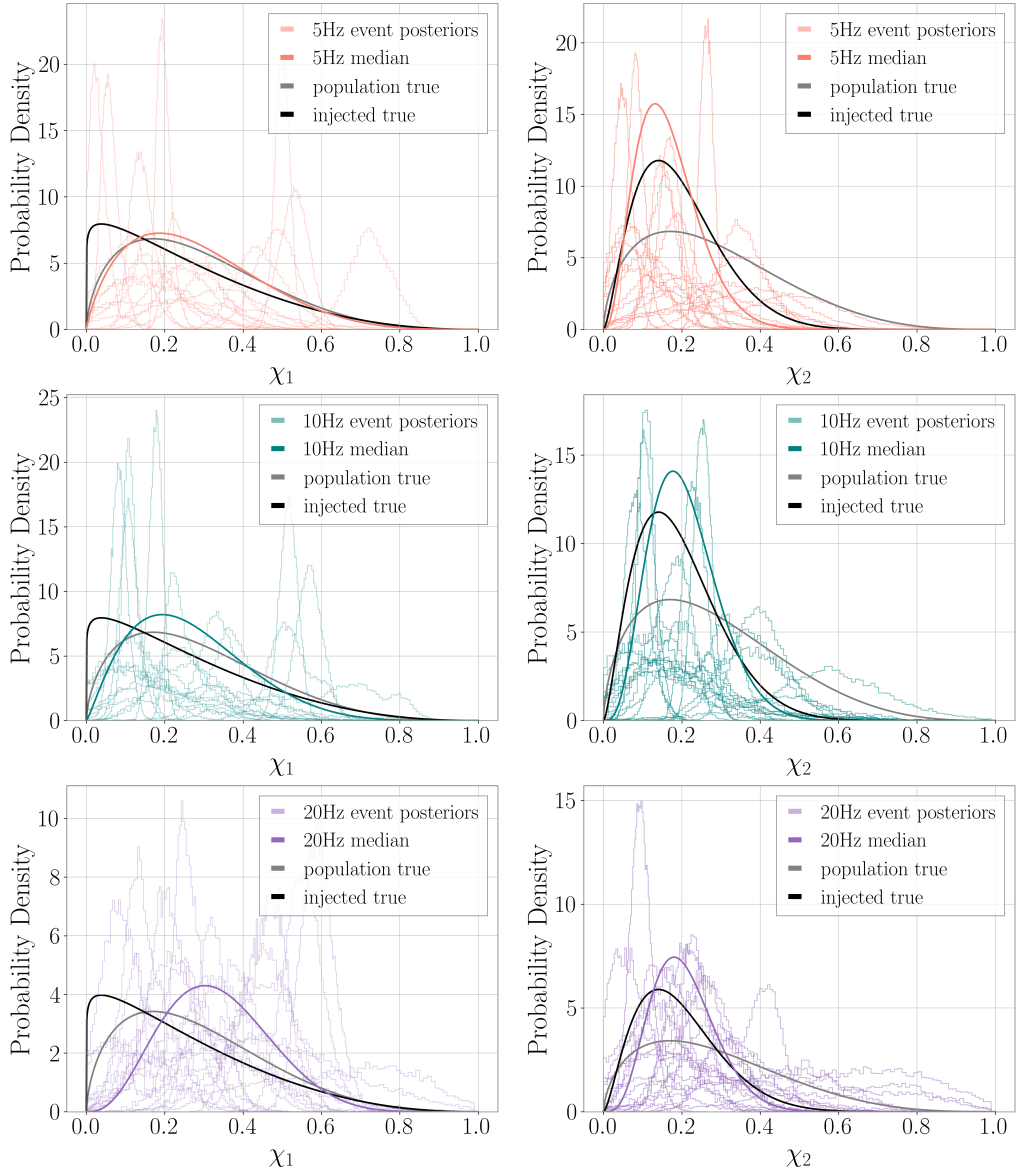


Figure 6.16: Posteriors for primary spin magnitude χ_1 on the left, and secondary spin magnitude χ_2 on the right. The top row shows results with $f_{\text{low}} = 5\text{Hz}$, the middle row $f_{\text{low}} = 10\text{Hz}$, and the bottom row $f_{\text{low}} = 20\text{Hz}$. The fainter histograms show individual event posteriors of spin magnitudes for the 20 binaries, plotted in light orange for 5Hz, light teal for 10Hz, and light purple for 20Hz. We also show, plotted in thicker orange, teal and purple, the Beta distributions of the median recovered posterior samples of $\{\alpha_{\chi_1}, \beta_{\chi_1}, \alpha_{\chi_2}, \beta_{\chi_2}\}$ for each value of f_{low} . In grey, we show Beta distributions for the primary and spin magnitude with the underlying population hyperparameter values (column 3 of Tab. 6.4), and in black are distributions with the injected hyperparameter values for the population of 20 binaries (column 4 of Tab. 6.4).

the small number of binaries with high primary spin magnitudes whose posteriors narrow with decreasing f_{eff} . This has the effect of adding more posterior weight to higher χ_1 values relative to low ones, and so will force the recovered population distribution of χ_1 towards larger values relative to the true injected population distribution. Therefore, although the peak of the recovered median distribution does tend towards lower values as f_{low} decreases and the broad individual χ_1 posteriors become narrower, the median distribution at 5Hz still peaks at larger values of χ_1 than the true population of 20 binaries would suggest. Consequently as f_{low} decreases the recovered Beta distribution more closely resembles the distribution of the underlying population in grey than the injected population in black.

In contrast to this, we see on the right side of Fig. 6.16 that the peak of the median recovered Beta distribution for χ_2 does tend towards that of our injected population in black as f_{low} decreases. This is because, unlike for χ_1 , the individual event posteriors for events with low χ_2 values are localised better with decreasing f_{low} , and so the peak of the distribution moves towards lower values. We finally note that while the peak of the recovered median distribution at 5Hz closely matches that of the injected population, the width is significantly narrower. This is as a result of the unconstrained β_{χ_2} value which has significant support at high values, and further supports the hypothesis that there is a degeneracy between the shape parameters α_{χ_i} and β_{χ_i} .

6.4 Discussion

In this Chapter we have assessed the impact of a varying low frequency cutoff upon constraining precession in a population of BBHs observed by next-generation detector networks. We first introduced the formalism of hierarchical inference in Sec. 6.1, and outlined the specific population hypermodels we use in Subsecs. 6.1.1 and Subsec. 6.1.2. To disentangle the effects of increasing the SNR and a larger number of observed precession cycles when decreasing f_{low} , we first performed systematic series of injections for a GW190521-like binary in Sec. 6.2. We perform one series where as f_{low} was lowered,

the network SNR ρ_{network} increased due to the longer observed signal, and a second series where f_{low} was varied but d_L was also adjusted to fix the network SNR. We observed that, as expected, as we lower the minimum cutoff frequency and the resulting signal SNR increases, this leads to reduced uncertainty on the recovered binary parameters. Specifically, we found that our results for the spins with $f_{\text{low}} = 20\text{Hz}$ were prior-driven, whereas the analyses with 10Hz and 5Hz more effectively recovered the injected binary parameters. Additionally, we also showed in Fig. 6.3 that if we fix the SNR while varying f_{low} , we also see decreasing uncertainty in the recovered posteriors due to the increased length of signal and increased number of observed precession cycles. We may therefore conclude that with better low frequency content there are two complementary effects at play that reduce the uncertainty in the inferred parameters: the increased SNR due to the longer signal, and the information gained from observing more of the binary evolution. We also note that for very high mass events such as GW190521, we observe only the merger-ringdown in-band with current detectors. Lowering the minimum frequency to 10Hz or 5Hz therefore brings the inspiral into sensitivity, allowing important constraints to be made on the population of IMBH binaries and for us to explore the upper mass gap.

We then assessed the impact of a lower frequency cutoff on the inference of the astrophysical population of binaries in Sec. 6.3. To do so, we drew a representative sample of 20 binaries from the population inferred from GWTC-3, but then noted in Subsec. 6.3.1 that due to small number statistics, our population of 20 binaries will have a different distribution than the true underlying population, and so different values for the hyperparameters for each population hypermodel. We fit our sample of 20 binaries to each of the hypermodel distributions to obtain more representative injected values for the hyperparameters, but noted that a bivariate Gaussian in the effective spin parameters is not a good fit to either the underlying population distribution or our subset. We then showed our results for the individual binary parameter estimation in Subsec. 6.3.2 highlighting the mass and spin results for two specific binaries, and saw that in general, a lower value of f_{low} leads to better recovered spin parameters with smaller uncertainty, and that this

effect is particularly pronounced for the effective spin parameters and especially χ_{eff} .

We then used this sample of 20 binaries to perform population inference for two different spin hypermodels in Sec. 6.3. We presented our results for the **Gaussian Spin** hypermodel in Subsec. 6.3.3, and for the **Beta Magnitudes** hypermodel in Subsubsec. 6.3.3. For the **Gaussian Spin** hypermodel, we found that a lower frequency cutoff led to a better localisation of the injected population hyperparameters, as the bias between the median recovered values and the true values decreased with decreasing f_{low} . We also saw that while the 20Hz analysis appeared to better recover the hyperparameters of the underlying population of binaries from which we drew our sample of 20, these results were actually prior-driven. However, as f_{low} decreased the posteriors became constrained away from the prior and better matched the injected hyperparameter values. We note, however, that since the bivariate Gaussian on χ_p and χ_{eff} is not a good fit to either the underlying population distribution nor the distribution of our 20 binaries, this hypermodel has its limitations and the results may not accurately reflect the potential improvement to be gained in constraining the effective spin parameters at lower f_{low} . We leave analysis using different effective spin hypermodels to future work.

For the **Beta Magnitudes** hypermodel, we found that the dominant effect of lowering the cutoff frequency was to reduce the bias between the recovered hyperparameter posteriors and the true injected value for the population of 20 binaries we considered, for three out of the four hyperparameters which make up this hypermodel. We also found for the primary spin magnitude χ_1 that a lower f_{low} led to tighter constraints on the shape of the population distribution, as evidenced by narrowing 90% credible interval widths for the hyperparameters α_{χ_1} and β_{χ_1} . We found that our analyses consistently overestimated the values of all the hyperparameters for this hypermodel, meaning that the recovered spin distributions are narrower than those injected.

As we only considered a population of 20 binaries in this work for reasons of computational cost, we are limited by small number statistics in our analyses. We have quantified this effect by calculating the different hyperparameter values of our injected population as

well as the underlying population distribution we drew from, and choosing these injected values as our true values to recover. However, we note that we have not taken selection effects into account in our Bayesian hierarchical analyses [233–235], which may go some way towards helping recover the underlying population distribution. Next-generation ground-based detectors like ET or CE are expected to observe $\sim 10^5 - 10^6$ BBH mergers per year [253], and so these small number statistics will not come into play as we have seen here, but selection effects will need to be carefully considered to ensure unbiased population inference.

We also note that we have used the same 20 binaries in the analysis for each value of f_{low} . In a more realistic scenario, the number of detected BBH mergers would increase with better low frequency sensitivity, as more events pass the SNR threshold for detection. A larger population with lower f_{low} would feasibly result in better constraints on the population spin distributions, leading to narrower distributions on the hyperparameters. This may explain why the dominant effect we have seen in our results is that the bias between the true injected value and the median posterior values decreases with f_{low} , while the posterior widths remains approximately constant: We are more accurately recovering but not more precisely constraining the population spin distributions.

The next generation of ground-based GW detectors will facilitate a dramatic change in the number of BBH merger events we detect, as well as the accuracy with which we are able to recover their binary source parameters. Though the exact low frequency sensitivity of these instruments remains uncertain, we have shown that in order to take full advantage of these detections to accurately constrain precession and spin distributions of the population, a low frequency cutoff lower than 20Hz may be critical, with further improved prospects for population measurements with f_{low} as low as 5Hz.

7. Conclusion

No confident detection of precession in a BBH merger has yet been made. Despite promising candidate events such as GW190521 [37] and GW200129 [38], inference of precession is obfuscated by the short length of the GW190521 signal due to its high mass [37], and detection of precession in GW200129 is greatly dependent upon data quality considerations [217] and precessing waveform systematics [218]. This lack of a confident detection is in contrast with measurements of precession on a population level [13, 34]. Consequently, with the fourth LVK observing run just beginning we look towards future observations of unambiguously precessing binaries, both with current instruments and in next-generation detector networks.

The extraction of astrophysical information from BBH merger GW signals is predicated on the availability of waveform models. In order to observe and constrain astrophysical phenomena such as spin precession and higher modes, waveform models need to have these phenomena built into them. Additionally, waveform systematics can have significant impact on inference of precession, for example as was seen with GW200129 [218]. Therefore calibrating semianalytic waveform models to precessing NR including higher modes to ensure their accuracy is crucial. In Chap. 4, we developed and tested a precessing spin dimensional reduction that faithfully matches precessing higher-mode waveforms in the strong-field regime with fewer spin degrees of freedom. This effective precession spin provides a promising pathway towards meaningful calibration of multimodal precessing semianalytic waveform models to precessing NR, which could reduce precessing waveform systematics for future detections.

Completeness and accuracy of waveform models can come at the cost of computational efficiency. A typical PE run for a BBH comprises around $10^6 - 10^8$ waveform evaluations, so we need to ensure our models efficient enough for practical use. Surrogate modelling is a technique used to build models which are faithful to an underlying (accurate and physically complete) waveform model, but can be evaluated much more efficiently. In Chap. 5 we

built and tested a surrogate model of the precessing higher mode EOB model `SEOBNRv4PHM` [48]. Our model `SEOBNN_v4PHM_4dq2` is accurate to the underlying model over its training parameter space range, and utilises artificial neural networks to achieve speeds two orders of magnitude faster than `SEOBNRv4PHM`. The neural network implementation also allows for leverage of GPUs, to achieve a further 50 times speedup. If we take waveform evaluations to be the dominant cost of a PE analysis, this significant speed up is equivalent to reducing run time from 2 weeks to just 4 minutes.

Looking towards the future of GW science, we anticipate a network of next-generation ground-based detectors with greatly improved sensitivity compared to current instruments [225, 226]. We have seen with GW190521 that very heavy binaries with extremely short signals leave room for many alternative interpretations to a quasicircular precessing BBH merger [210–216], and so improved low frequency sensitivity of next-generation detectors could have a significant impact upon constraining precession, both for individual events and across the binary population. In Chap. 6 we assessed the impact of different minimum cutoff frequencies f_{low} upon our ability to constrain precession for individual binaries and across a population of 20 BBHs. We found that a lower frequency of 5Hz or even 10Hz as opposed to 20Hz allows us to more accurately measure the population spin distributions, even with a relatively small population. Next-generation detectors are expected to have improved low frequency sensitivity through cryogenic technology and reduced seismic noise, though the exact minimum frequency cutoff is still uncertain. We have shown that improving the sensitivity below 20Hz, will have a significant impact upon our ability to measure the true BBH population spin distributions and subsequently constrain formation channels of BBHs [21–29, 31–36].

Bibliography

- [1] Loïc Lannelongue, Jason Grealey, and Michael Inouye. “Green Algorithms: Quantifying the Carbon Footprint of Computation”. *Adv. Sci.* 8.12 (2021). DOI: [10.1002/advs.202100707](https://doi.org/10.1002/advs.202100707).
- [2] Albert Einstein. “The foundation of the general theory of relativity.” *Annalen Phys.* 49.7 (1916). Ed. by Jong-Ping Hsu and D. Fine, pp. 769–822. DOI: [10.1002/andp.19163540702](https://doi.org/10.1002/andp.19163540702).
- [3] Alexander H. Nitz et al. “1-OGC: The First Open Gravitational-wave Catalog of Binary Mergers from Analysis of Public Advanced LIGO Data”. *The Astrophysical Journal* 872.2 (2019), p. 195. ISSN: 0004-637X. DOI: [10.3847/1538-4357/ab0108](https://doi.org/10.3847/1538-4357/ab0108). arXiv: [1811.01921](https://arxiv.org/abs/1811.01921) [[gr-qc](#)].
- [4] Alexander H. Nitz et al. “2-OGC: Open Gravitational-wave Catalog of Binary Mergers from Analysis of Public Advanced LIGO and Virgo Data”. *Astrophys. J.* 891 (2019), p. 123. DOI: [10.3847/1538-4357/ab733f](https://doi.org/10.3847/1538-4357/ab733f). arXiv: [1910.05331](https://arxiv.org/abs/1910.05331) [[astro-ph.HE](#)].
- [5] Alexander H. Nitz et al. “3-OGC: Catalog of Gravitational Waves from Compact-binary Mergers”. *The Astrophysical Journal* 922.1 (2021), p. 76. ISSN: 0004-637X. DOI: [10.3847/1538-4357/ac1c03](https://doi.org/10.3847/1538-4357/ac1c03). arXiv: [2105.09151](https://arxiv.org/abs/2105.09151) [[astro-ph.HE](#)].
- [6] R. Abbott et al. “Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences”. *Astrophys. J. Lett.* 915.1 (2021), p. L5. DOI: [10.3847/2041-8213/ac082e](https://doi.org/10.3847/2041-8213/ac082e). arXiv: [2106.15163](https://arxiv.org/abs/2106.15163) [[astro-ph.HE](#)].

- [7] Tejaswi Venumadhav et al. “New Binary Black Hole Mergers in the Second Observing Run of Advanced LIGO and Advanced Virgo”. *Phys. Rev. D* 101.8 (2020), p. 083030. DOI: [10.1103/PhysRevD.101.083030](https://doi.org/10.1103/PhysRevD.101.083030). arXiv: [1904.07214](https://arxiv.org/abs/1904.07214) [[astro-ph.HE](#)].
- [8] B. P. Abbott et al. “Tests of general relativity with GW150914”. *Phys. Rev. Lett.* 116.22 (2016). [Erratum: *Phys.Rev.Lett.* 121, 129902 (2018)], p. 221101. DOI: [10.1103/PhysRevLett.116.221101](https://doi.org/10.1103/PhysRevLett.116.221101). arXiv: [1602.03841](https://arxiv.org/abs/1602.03841) [[gr-qc](#)].
- [9] B. P. Abbott et al. “Tests of General Relativity with GW170817”. *Phys. Rev. Lett.* 123.1 (2019), p. 011102. DOI: [10.1103/PhysRevLett.123.011102](https://doi.org/10.1103/PhysRevLett.123.011102). arXiv: [1811.00364](https://arxiv.org/abs/1811.00364) [[gr-qc](#)].
- [10] B. P. Abbott et al. “Tests of General Relativity with the Binary Black Hole Signals from the LIGO-Virgo Catalog GWTC-1”. *Phys. Rev. D* 100.10 (2019), p. 104036. DOI: [10.1103/PhysRevD.100.104036](https://doi.org/10.1103/PhysRevD.100.104036). arXiv: [1903.04467](https://arxiv.org/abs/1903.04467) [[gr-qc](#)].
- [11] R. Abbott et al. “Tests of general relativity with binary black holes from the second LIGO-Virgo gravitational-wave transient catalog”. *Phys. Rev. D* 103.12 (2021), p. 122002. DOI: [10.1103/PhysRevD.103.122002](https://doi.org/10.1103/PhysRevD.103.122002). arXiv: [2010.14529](https://arxiv.org/abs/2010.14529) [[gr-qc](#)].
- [12] R. Abbott et al. “Tests of General Relativity with GWTC-3” (2021). arXiv: [2112.06861](https://arxiv.org/abs/2112.06861) [[gr-qc](#)].
- [13] Virgo Collaboration LIGO Scientific Collaboration and KAGRA Collaboration et al. “Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3”. *Physical Review X* 13.1 (2023), p. 011048. DOI: [10.1103/PhysRevX.13.011048](https://doi.org/10.1103/PhysRevX.13.011048). arXiv: [2111.03634](https://arxiv.org/abs/2111.03634) [[astro-ph.HE](#)].
- [14] Theocharis A. Apostolatos et al. “Spin Induced Orbital Precession and Its Modulation of the Gravitational Wave Forms from Merging Binaries”. *Phys. Rev. D* 49 (1994), pp. 6274–6297. DOI: [10.1103/PhysRevD.49.6274](https://doi.org/10.1103/PhysRevD.49.6274).
- [15] Lawrence E. Kidder. “Coalescing Binary Systems of Compact Objects to post-Newtonian 5/2 Order. 5. Spin Effects”. *Phys. Rev. D* 52 (1995), pp. 821–847. DOI: [10.1103/PhysRevD.52.821](https://doi.org/10.1103/PhysRevD.52.821). arXiv: [gr-qc/9506022](https://arxiv.org/abs/gr-qc/9506022).

- [16] M.V. van der Sluys et al. “Gravitational-Wave Astronomy with Inspiral Signals of Spinning Compact-Object Binaries”. *Astrophys. J. Lett.* 688 (2008), p. L61. DOI: [10.1086/595279](https://doi.org/10.1086/595279). arXiv: [0710.1897](https://arxiv.org/abs/0710.1897) [[astro-ph](#)].
- [17] Hee-Suk Cho et al. “Gravitational Waves from Black Hole-Neutron Star Binaries: Effective Fisher Matrices and Parameter Estimation Using Higher Harmonics”. *Phys. Rev. D* 87.2 (2013), p. 024004. DOI: [10.1103/PhysRevD.87.024004](https://doi.org/10.1103/PhysRevD.87.024004). arXiv: [1209.4494](https://arxiv.org/abs/1209.4494) [[gr-qc](#)].
- [18] R. O’Shaughnessy et al. “Parameter Estimation of Gravitational Waves from Precessing Black Hole-Neutron Star Inspirals with Higher Harmonics”. *Phys. Rev. D* 89.10 (2014), p. 102005. DOI: [10.1103/PhysRevD.89.102005](https://doi.org/10.1103/PhysRevD.89.102005). arXiv: [1403.0544](https://arxiv.org/abs/1403.0544) [[gr-qc](#)].
- [19] Geraint Pratten et al. “Measuring precession in asymmetric compact binaries”. *Phys. Rev. Res.* 2.4 (2020), p. 043096. DOI: [10.1103/PhysRevResearch.2.043096](https://doi.org/10.1103/PhysRevResearch.2.043096). arXiv: [2006.16153](https://arxiv.org/abs/2006.16153) [[gr-qc](#)].
- [20] Vassiliki Kalogera. “Spin orbit misalignment in close binaries with two compact objects”. *Astrophys. J.* 541 (2000), pp. 319–328. DOI: [10.1086/309400](https://doi.org/10.1086/309400). arXiv: [astro-ph/9911417](https://arxiv.org/abs/astro-ph/9911417).
- [21] Ilya Mandel and Richard O’Shaughnessy. “Compact Binary Coalescences in the Band of Ground-based Gravitational-Wave Detectors”. *Class. Quant. Grav.* 27 (2010). Ed. by Sascha Husa and Badri Krishnan, p. 114007. DOI: [10.1088/0264-9381/27/11/114007](https://doi.org/10.1088/0264-9381/27/11/114007). arXiv: [0912.1074](https://arxiv.org/abs/0912.1074) [[astro-ph.HE](#)].
- [22] Carl L. Rodriguez et al. “Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO”. *Phys. Rev. Lett.* 115.5 (2015). [Erratum: *Phys.Rev.Lett.* 116, 029901 (2016)], p. 051101. DOI: [10.1103/PhysRevLett.115.051101](https://doi.org/10.1103/PhysRevLett.115.051101). arXiv: [1505.00792](https://arxiv.org/abs/1505.00792) [[astro-ph.HE](#)].

- [23] Carl L. Rodriguez et al. “Illuminating Black Hole Binary Formation Channels with Spins in Advanced LIGO”. *Astrophys. J. Lett.* 832.1 (2016), p. L2. DOI: [10.3847/2041-8205/832/1/L2](https://doi.org/10.3847/2041-8205/832/1/L2). arXiv: [1609.05916](https://arxiv.org/abs/1609.05916) [[astro-ph.HE](#)].
- [24] Will M. Farr et al. “Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves”. *Nature* 548 (2017), p. 426. DOI: [10.1038/nature23453](https://doi.org/10.1038/nature23453). arXiv: [1706.01385](https://arxiv.org/abs/1706.01385) [[astro-ph.HE](#)].
- [25] Ben Farr, Daniel E. Holz, and Will M. Farr. “Using Spin to Understand the Formation of LIGO and Virgo’s Black Holes”. *Astrophys. J. Lett.* 854.1 (2018), p. L9. DOI: [10.3847/2041-8213/aaaa64](https://doi.org/10.3847/2041-8213/aaaa64). arXiv: [1709.07896](https://arxiv.org/abs/1709.07896) [[astro-ph.HE](#)].
- [26] Maya Fishbach, Daniel E. Holz, and Ben Farr. “Are LIGO’s Black Holes Made From Smaller Black Holes?” *Astrophys. J. Lett.* 840.2 (2017), p. L24. DOI: [10.3847/2041-8213/aa7045](https://doi.org/10.3847/2041-8213/aa7045). arXiv: [1703.06869](https://arxiv.org/abs/1703.06869) [[astro-ph.HE](#)].
- [27] K. Belczynski et al. “Evolutionary roads leading to low effective spins, high black hole masses, and O1/O2 rates for LIGO/Virgo binary black holes”. *Astron. Astrophys.* 636 (2020), A104. DOI: [10.1051/0004-6361/201936528](https://doi.org/10.1051/0004-6361/201936528). arXiv: [1706.07053](https://arxiv.org/abs/1706.07053) [[astro-ph.HE](#)].
- [28] Davide Gerosa and Emanuele Berti. “Are merging black holes born from stellar collapse or previous mergers?” *Phys. Rev. D* 95.12 (2017), p. 124046. DOI: [10.1103/PhysRevD.95.124046](https://doi.org/10.1103/PhysRevD.95.124046). arXiv: [1703.06223](https://arxiv.org/abs/1703.06223) [[gr-qc](#)].
- [29] Colm Talbot and Eric Thrane. “Determining the population properties of spinning black holes”. *Phys. Rev. D* 96.2 (2017), p. 023012. DOI: [10.1103/PhysRevD.96.023012](https://doi.org/10.1103/PhysRevD.96.023012). arXiv: [1704.08370](https://arxiv.org/abs/1704.08370) [[astro-ph.HE](#)].
- [30] Xingjiang Zhu et al. “Inferring the population properties of binary neutron stars with gravitational-wave measurements of spin”. *Phys. Rev. D* 98 (2018), p. 043002. DOI: [10.1103/PhysRevD.98.043002](https://doi.org/10.1103/PhysRevD.98.043002). arXiv: [1711.09226](https://arxiv.org/abs/1711.09226) [[astro-ph.HE](#)].

-
- [31] Simon Stevenson, Christopher P. L. Berry, and Ilya Mandel. “Hierarchical analysis of gravitational-wave measurements of binary black hole spin–orbit misalignments”. *Mon. Not. Roy. Astron. Soc.* 471.3 (2017), pp. 2801–2811. DOI: [10.1093/mnras/stx1764](https://doi.org/10.1093/mnras/stx1764). arXiv: [1703.06873](https://arxiv.org/abs/1703.06873) [[astro-ph.HE](#)].
- [32] Daniel Wysocki, Jacob Lange, and Richard O’Shaughnessy. “Reconstructing phenomenological distributions of compact binaries via gravitational wave observations”. *Phys. Rev. D* 100.4 (2019), p. 043012. DOI: [10.1103/PhysRevD.100.043012](https://doi.org/10.1103/PhysRevD.100.043012). arXiv: [1805.06442](https://arxiv.org/abs/1805.06442) [[gr-qc](#)].
- [33] 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 (2019), p. L24. DOI: [10.3847/2041-8213/ab3800](https://doi.org/10.3847/2041-8213/ab3800). arXiv: [1811.12940](https://arxiv.org/abs/1811.12940) [[astro-ph.HE](#)].
- [34] R. Abbott et al. “Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog”. *Astrophys. J. Lett.* 913.1 (2021), p. L7. DOI: [10.3847/2041-8213/abe949](https://doi.org/10.3847/2041-8213/abe949). arXiv: [2010.14533](https://arxiv.org/abs/2010.14533) [[astro-ph.HE](#)].
- [35] R. Abbott et al. “Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3”. *Phys. Rev. X* 13.1 (2023), p. 011048. DOI: [10.1103/PhysRevX.13.011048](https://doi.org/10.1103/PhysRevX.13.011048). arXiv: [2111.03634](https://arxiv.org/abs/2111.03634) [[astro-ph.HE](#)].
- [36] Chase Kimball et al. “Black Hole Genealogy: Identifying Hierarchical Mergers with Gravitational Waves”. *Astrophys. J.* 900.2 (2020), p. 177. DOI: [10.3847/1538-4357/aba518](https://doi.org/10.3847/1538-4357/aba518). arXiv: [2005.00023](https://arxiv.org/abs/2005.00023) [[astro-ph.HE](#)].
- [37] Richard Abbott et al. “GW190521: A Binary Black Hole Merger with a Total Mass of \odot ”. *Phys. Rev. Lett.* 125.10 (2020), p. 101102. DOI: [10.1103/PhysRevLett.125.101102](https://doi.org/10.1103/PhysRevLett.125.101102). arXiv: [2009.01075](https://arxiv.org/abs/2009.01075) [[gr-qc](#)].
- [38] Mark Hannam et al. “General-Relativistic Precession in a Black-Hole Binary”. *Nature* 610.7933 (2022), pp. 652–655. ISSN: 1476-4687. DOI: [10.1038/s41586-022-05212-z](https://doi.org/10.1038/s41586-022-05212-z). arXiv: [2112.11300](https://arxiv.org/abs/2112.11300) [[gr-qc](#)].

- [39] Lucy M. Thomas, Patricia Schmidt, and Geraint Pratten. “New Effective Precession Spin for Modeling Multimodal Gravitational Waveforms in the Strong-Field Regime”. *Physical Review D* 103.8 (2021), p. 083022. DOI: [10.1103/PhysRevD.103.083022](https://doi.org/10.1103/PhysRevD.103.083022). arXiv: [2012.02209 \[gr-qc\]](https://arxiv.org/abs/2012.02209).
- [40] Lucy M. Thomas, Geraint Pratten, and Patricia Schmidt. “Accelerating Multimodal Gravitational Waveforms from Precessing Compact Binaries with Artificial Neural Networks”. *Physical Review D* 106.10 (2022), p. 104029. DOI: [10.1103/PhysRevD.106.104029](https://doi.org/10.1103/PhysRevD.106.104029). arXiv: [2205.14066 \[gr-qc\]](https://arxiv.org/abs/2205.14066).
- [41] Éanna É Flanagan and Scott A. Hughes. “The Basics of Gravitational Wave Theory”. *New Journal of Physics* 7.1 (2005), p. 204. ISSN: 1367-2630. DOI: [10.1088/1367-2630/7/1/204](https://doi.org/10.1088/1367-2630/7/1/204). arXiv: [gr-qc/0501041](https://arxiv.org/abs/gr-qc/0501041).
- [42] Steven Weinberg. *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. New York: John Wiley and Sons, 1972. ISBN: 978-0-471-92567-5.
- [43] Sean M. Carroll. *Spacetime and Geometry*. Cambridge University Press, 2019. ISBN: 978-0-8053-8732-2.
- [44] Curt Cutler and Éanna E. Flanagan. “Gravitational Waves from Merging Compact Binaries: How Accurately Can One Extract the Binary’s Parameters from the Inspiral Waveform?” *Physical Review D* 49.6 (1994), pp. 2658–2697. DOI: [10.1103/PhysRevD.49.2658](https://doi.org/10.1103/PhysRevD.49.2658). arXiv: [gr-qc/9402014](https://arxiv.org/abs/gr-qc/9402014).
- [45] Richard 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” (2021). arXiv: [2108.01045 \[gr-qc\]](https://arxiv.org/abs/2108.01045). URL: <https://arxiv.org/abs/2108.01045>.
- [46] Vassiliki Kalogera and Gordon Baym. “The Maximum Mass of a Neutron Star”. *Astrophys. J. Lett.* 470 (1996), pp. L61–L64. DOI: [10.1086/310296](https://doi.org/10.1086/310296). arXiv: [astro-ph/9608059](https://arxiv.org/abs/astro-ph/9608059).

- [47] Saul A. Teukolsky. “Perturbations of a Rotating Black Hole. 1. Fundamental Equations for Gravitational Electromagnetic and Neutrino Field Perturbations”. *Astrophys. J.* 185 (1973), pp. 635–647. DOI: [10.1086/152444](https://doi.org/10.1086/152444).
- [48] Serguei Ossokine et al. “Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation”. *Phys. Rev. D* 102.4 (2020), p. 044055. DOI: [10.1103/PhysRevD.102.044055](https://doi.org/10.1103/PhysRevD.102.044055). arXiv: [2004.09442](https://arxiv.org/abs/2004.09442) [gr-qc].
- [49] LIGO Scientific Collaboration. *LIGO Algorithm Library - LALSuite*. 2018. DOI: [10.7935/GT1W-FZ16](https://doi.org/10.7935/GT1W-FZ16).
- [50] Benjamin P Abbott et al. “A Guide to LIGO–Virgo Detector Noise and Extraction of Transient Gravitational-Wave Signals”. *Class. Quant. Grav.* 37.5 (2020), p. 055002. DOI: [10.1088/1361-6382/ab685e](https://doi.org/10.1088/1361-6382/ab685e). arXiv: [1908.11170](https://arxiv.org/abs/1908.11170) [gr-qc].
- [51] LIGO Scientific Collaboration and The Virgo Collaboration. *Noise Curves Used for Simulations in the Update of the Observing Scenarios Paper*. 2019. URL: <https://dcc.ligo.org/LIGO-T2000012/public>.
- [52] Matthew Evans et al. *Cosmic Explorer Strain Sensitivity*. 2021. URL: <https://dcc.cosmicexplorer.org/cgi-bin/DocDB/ShowDocument?.submit=Identifier&docid=T2000017&version=>.
- [53] Michele Punturo and Einstein Telescope Scientific Collaboration. *Einstein Telescope Sensitivities*. 2021. URL: <https://www.et-gw.eu/index.php/etsensitivities>.
- [54] P. C. Peters. “Gravitational Radiation and the Motion of Two Point Masses”. *Physical Review* 136.4B (1964), B1224–B1232. DOI: [10.1103/PhysRev.136.B1224](https://doi.org/10.1103/PhysRev.136.B1224).
- [55] B. S. Sathyaprakash and B. F. Schutz. “Physics, Astrophysics and Cosmology with Gravitational Waves”. *Living Reviews in Relativity* 12.1 (2009), p. 2. ISSN: 2367-3613, 1433-8351. DOI: [10.12942/lrr-2009-2](https://doi.org/10.12942/lrr-2009-2). arXiv: [0903.0338](https://arxiv.org/abs/0903.0338) [gr-qc].

- [56] R. Abbott, T.D. Abbott, S. Abraham, et al. “Diving below the Spin-down Limit: Constraints on Gravitational Waves from the Energetic Young Pulsar PSR J0537-6910”. *The Astrophysical Journal Letters* 913.2 (2021), p. L27. ISSN: 2041-8205, 2041-8213. DOI: [10.3847/2041-8213/abffcd](https://doi.org/10.3847/2041-8213/abffcd). arXiv: [2012.12926](https://arxiv.org/abs/2012.12926) [[astro-ph.HE](#)].
- [57] Luc Blanchet. “Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries”. *Living Reviews in Relativity* 17.1 (2014), p. 2. ISSN: 1433-8351. DOI: [10.12942/lrr-2014-2](https://doi.org/10.12942/lrr-2014-2). arXiv: [1310.1528](https://arxiv.org/abs/1310.1528) [[gr-qc](#)].
- [58] Alessandra Buonanno et al. “Comparison of Post-Newtonian Templates for Compact Binary Inspiral Signals in Gravitational-Wave Detectors”. *Physical Review D* 80.8 (2009), p. 084043. DOI: [10.1103/PhysRevD.80.084043](https://doi.org/10.1103/PhysRevD.80.084043). arXiv: [0907.0700](https://arxiv.org/abs/0907.0700) [[gr-qc](#)].
- [59] Lawrence E. Kidder. “Using Full Information When Computing Modes of Post-Newtonian Waveforms from Inspiralling Compact Binaries in Circular Orbit”. *Physical Review D* 77.4 (2008), p. 044016. DOI: [10.1103/PhysRevD.77.044016](https://doi.org/10.1103/PhysRevD.77.044016). arXiv: [0710.0614](https://arxiv.org/abs/0710.0614) [[gr-qc](#)].
- [60] Kip S. Thorne. “Multipole Expansions of Gravitational Radiation”. *Reviews of Modern Physics* 52.2 (1980), pp. 299–339. DOI: [10.1103/RevModPhys.52.299](https://doi.org/10.1103/RevModPhys.52.299).
- [61] M. Alcubierre. *Introduction to 3+1 Numerical Relativity*. International Series of Monographs on Physics. OUP Oxford, 2008. ISBN: 978-0-19-154829-1. URL: <https://books.google.co.uk/books?id=-xRREAAAQBAJ>.
- [62] Thomas W. Baumgarte and Stuart L. Shapiro. *Numerical Relativity: Solving Einstein’s Equations on the Computer*. Cambridge University Press, 2010. DOI: [10.1017/CBO9781139193344](https://doi.org/10.1017/CBO9781139193344).
- [63] Richard Arnowitt, Stanley Deser, and Charles W. Misner. “Republication of: The Dynamics of General Relativity”. *General Relativity and Gravitation* 40.9 (2008), pp. 1997–2027. ISSN: 1572-9532. DOI: [10.1007/s10714-008-0661-1](https://doi.org/10.1007/s10714-008-0661-1). arXiv: [gr-qc/0405109](https://arxiv.org/abs/gr-qc/0405109).

-
- [64] Thomas W. Baumgarte and Stuart L. Shapiro. “Numerical Integration of Einstein’s Field Equations”. *Physical Review D* 59.2 (1998), p. 024007. DOI: [10.1103/PhysRevD.59.024007](https://doi.org/10.1103/PhysRevD.59.024007). arXiv: [gr-qc/9810065](https://arxiv.org/abs/gr-qc/9810065).
- [65] Masaru Shibata and Takashi Nakamura. “Evolution of Three-Dimensional Gravitational Waves: Harmonic Slicing Case”. *Physical Review D* 52.10 (1995), pp. 5428–5444. DOI: [10.1103/PhysRevD.52.5428](https://doi.org/10.1103/PhysRevD.52.5428).
- [66] Frans Pretorius. “Evolution of Binary Black-Hole Spacetimes”. *Physical Review Letters* 95.12 (2005), p. 121101. DOI: [10.1103/PhysRevLett.95.121101](https://doi.org/10.1103/PhysRevLett.95.121101). arXiv: [gr-qc/0507014](https://arxiv.org/abs/gr-qc/0507014).
- [67] J. Thornburg. “Coordinates and Boundary Conditions for the General Relativistic Initial Data Problem”. *Classical and Quantum Gravity* 4.5 (1987), p. 1119. ISSN: 0264-9381. DOI: [10.1088/0264-9381/4/5/013](https://doi.org/10.1088/0264-9381/4/5/013).
- [68] Eleanor Hamilton et al. “BAM Catalogue of Precessing Binary Black Hole Simulations” (2023).
- [69] Michael Boyle et al. “The SXS Collaboration Catalog of Binary Black Hole Simulations”. *Classical and Quantum Gravity* 36.19 (2019), p. 195006. ISSN: 0264-9381. DOI: [10.1088/1361-6382/ab34e2](https://doi.org/10.1088/1361-6382/ab34e2). arXiv: [1904.04831](https://arxiv.org/abs/1904.04831) [[gr-qc](#)].
- [70] Michela Mapelli. “Formation Channels of Single and Binary Stellar-Mass Black Holes”. 2021. DOI: [10.1007/978-981-15-4702-7_16-1](https://doi.org/10.1007/978-981-15-4702-7_16-1). arXiv: [2106.00699](https://arxiv.org/abs/2106.00699) [[astro-ph.HE](#)].
- [71] P. Ajith et al. “Inspirals-Merger-Ringdown Waveforms for Black-Hole Binaries with Nonprecessing Spins”. *Physical Review Letters* 106.24 (2011), p. 241101. DOI: [10.1103/PhysRevLett.106.241101](https://doi.org/10.1103/PhysRevLett.106.241101). arXiv: [0909.2867](https://arxiv.org/abs/0909.2867) [[gr-qc](#)].
- [72] Laura Santamaria et al. “Matching Post-Newtonian and Numerical Relativity Waveforms: Systematic Errors and a New Phenomenological Model for Non-Precessing Black Hole Binaries”. *Phys. Rev. D* 82 (2010), p. 064016. DOI: [10.1103/PhysRevD.82.064016](https://doi.org/10.1103/PhysRevD.82.064016). arXiv: [1005.3306](https://arxiv.org/abs/1005.3306) [[gr-qc](#)].

- [73] P. Ajith. “Addressing the Spin Question in Gravitational-Wave Searches: Waveform Templates for Inspiralling Compact Binaries with Nonprecessing Spins”. *Physical Review D* 84.8 (2011), p. 084037. DOI: [10.1103/PhysRevD.84.084037](https://doi.org/10.1103/PhysRevD.84.084037). arXiv: [1107.1267 \[gr-qc\]](https://arxiv.org/abs/1107.1267).
- [74] Christian Reisswig et al. “Gravitational-Wave Detectability of Equal-Mass Black-Hole Binaries with Aligned Spins”. *Physical Review D* 80.12 (2009), p. 124026. DOI: [10.1103/PhysRevD.80.124026](https://doi.org/10.1103/PhysRevD.80.124026). arXiv: [0907.0462 \[gr-qc\]](https://arxiv.org/abs/0907.0462).
- [75] Michael Pürrer et al. “Testing the Validity of the Single-Spin Approximation in Inspiral-Merger-Ringdown Waveforms”. *Physical Review D* 88.6 (2013), p. 064007. DOI: [10.1103/PhysRevD.88.064007](https://doi.org/10.1103/PhysRevD.88.064007). arXiv: [1306.2320 \[gr-qc\]](https://arxiv.org/abs/1306.2320).
- [76] Étienne Racine. “Analysis of Spin Precession in Binary Black Hole Systems Including Quadrupole-Monopole Interaction”. *Physical Review D* 78.4 (2008), p. 044021. DOI: [10.1103/PhysRevD.78.044021](https://doi.org/10.1103/PhysRevD.78.044021). arXiv: [0803.1820 \[gr-qc\]](https://arxiv.org/abs/0803.1820).
- [77] Kip S. Thorne and James B. Hartle. “Laws of Motion and Precession for Black Holes and Other Bodies”. *Physical Review D* 31.8 (1984), pp. 1815–1837. DOI: [10.1103/PhysRevD.31.1815](https://doi.org/10.1103/PhysRevD.31.1815).
- [78] B. M. Barker and R. F. O’Connell. “Gravitational Two-Body Problem with Arbitrary Masses, Spins, and Quadrupole Moments”. *Physical Review D* 12.2 (1975), pp. 329–335. DOI: [10.1103/PhysRevD.12.329](https://doi.org/10.1103/PhysRevD.12.329).
- [79] Alberto Vecchio. “LISA Observations of Rapidly Spinning Massive Black Hole Binary Systems”. *Phys. Rev. D* 70 (2004), p. 042001. DOI: [10.1103/PhysRevD.70.042001](https://doi.org/10.1103/PhysRevD.70.042001). arXiv: [astro-ph/0304051](https://arxiv.org/abs/astro-ph/0304051).
- [80] Ryan N. Lang and Scott A. Hughes. “Measuring Coalescing Massive Binary Black Holes with Gravitational Waves: The Impact of Spin-Induced Precession”. *Phys. Rev. D* 74 (2006), p. 122001. DOI: [10.1103/PhysRevD.75.089902](https://doi.org/10.1103/PhysRevD.75.089902). arXiv: [gr-qc/0608062](https://arxiv.org/abs/gr-qc/0608062).

-
- [81] Katerina Chatziioannou et al. “Spin-Precession: Breaking the Black Hole–Neutron Star Degeneracy”. *Astrophys. J. Lett.* 798.1 (2014), p. L17. DOI: [10.1088/2041-8205/798/1/L17](https://doi.org/10.1088/2041-8205/798/1/L17). arXiv: [1402.3581](https://arxiv.org/abs/1402.3581) [[gr-qc](#)].
- [82] 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 (2019), e010. ISSN: 1323-3580, 1448-6083. DOI: [10.1017/pasa.2019.2](https://doi.org/10.1017/pasa.2019.2). arXiv: [1809.02293](https://arxiv.org/abs/1809.02293) [[astro-ph.IM](#)].
- [83] John Skilling. “Nested Sampling”. *AIP Conference Proceedings* 735.1 (2004), pp. 395–405. ISSN: 0094-243X. DOI: [10.1063/1.1835238](https://doi.org/10.1063/1.1835238). eprint: https://pubs.aip.org/aip/acp/article-pdf/735/1/395/11702789/395\textbackslash_1\textbackslash_online.pdf.
- [84] N. Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. 21 (1953), pp. 1087–1092.
- [85] W. K. Hastings. “Monte Carlo Sampling Methods Using Markov Chains and Their Applications”. 57.1 (1970), pp. 97–109. ISSN: 00063444. JSTOR: [2334940](https://www.jstor.org/stable/2334940). URL: <http://www.jstor.org/stable/2334940>.
- [86] Joshua S. Speagle. “dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences”. *Mon. Not. Roy. Astron. Soc.* 493.3 (2020), pp. 3132–3158. DOI: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278). arXiv: [1904.02180](https://arxiv.org/abs/1904.02180) [[astro-ph.IM](#)].
- [87] Gregory Ashton et al. “Bilby: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy”. *The Astrophysical Journal Supplement Series* 241.2 (2019), p. 27. ISSN: 0067-0049. DOI: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc). arXiv: [1811.02042](https://arxiv.org/abs/1811.02042) [[astro-ph.IM](#)].
- [88] Alessandra Buonanno and Thibault Damour. “Effective One-Body Approach to General Relativistic Two-Body Dynamics”. *Phys. Rev. D* 59 (1999), p. 084006. DOI: [10.1103/PhysRevD.59.084006](https://doi.org/10.1103/PhysRevD.59.084006). arXiv: [gr-qc/9811091](https://arxiv.org/abs/gr-qc/9811091).

- [89] Jonathan Blackman et al. “Fast and Accurate Prediction of Numerical Relativity Waveforms from Binary Black Hole Coalescences Using Surrogate Models”. *Phys. Rev. Lett.* 115.12 (2015), p. 121102. DOI: [10.1103/PhysRevLett.115.121102](https://doi.org/10.1103/PhysRevLett.115.121102). arXiv: [1502.07758](https://arxiv.org/abs/1502.07758) [gr-qc].
- [90] Jonathan Blackman et al. “A Surrogate Model of Gravitational Waveforms from Numerical Relativity Simulations of Precessing Binary Black Hole Mergers”. *Phys. Rev. D* 95.10 (2017), p. 104023. DOI: [10.1103/PhysRevD.95.104023](https://doi.org/10.1103/PhysRevD.95.104023). arXiv: [1701.00550](https://arxiv.org/abs/1701.00550) [gr-qc].
- [91] Jonathan Blackman et al. “Numerical Relativity Waveform Surrogate Model for Generically Precessing Binary Black Hole Mergers”. *Phys. Rev. D* 96.2 (2017), p. 024058. DOI: [10.1103/PhysRevD.96.024058](https://doi.org/10.1103/PhysRevD.96.024058). arXiv: [1705.07089](https://arxiv.org/abs/1705.07089) [gr-qc].
- [92] Jun John Sakurai and Jim Napolitano. *Modern Quantum Mechanics*. Quantum Physics, Quantum Information and Quantum Computation. Cambridge University Press, 2020. ISBN: 978-0-8053-8291-4. DOI: [10.1017/9781108587280](https://doi.org/10.1017/9781108587280).
- [93] Ian Jones. *LIGO-T1200476-v3: Calculating Gravitational Waveforms: Examples*. URL: <https://dcc.ligo.org/LIGO-T1200476/public>.
- [94] Cameron Mills and Stephen Fairhurst. “Measuring Gravitational-Wave Higher-Order Multipoles”. *Physical Review D* 103.2 (2021), p. 024042. DOI: [10.1103/PhysRevD.103.024042](https://doi.org/10.1103/PhysRevD.103.024042). arXiv: [2007.04313](https://arxiv.org/abs/2007.04313) [gr-qc].
- [95] Vijay Varma et al. “Gravitational-Wave Observations of Binary Black Holes: Effect of Nonquadrupole Modes”. *Physical Review D* 90.12 (2014), p. 124004. DOI: [10.1103/PhysRevD.90.124004](https://doi.org/10.1103/PhysRevD.90.124004). arXiv: [1409.2349](https://arxiv.org/abs/1409.2349) [gr-qc].
- [96] Vijay Varma and Parameswaran Ajith. “Effects of Nonquadrupole Modes in the Detection and Parameter Estimation of Black Hole Binaries with Nonprecessing Spins”. *Physical Review D* 96.12 (2017), p. 124024. DOI: [10.1103/PhysRevD.96.124024](https://doi.org/10.1103/PhysRevD.96.124024). arXiv: [1612.05608](https://arxiv.org/abs/1612.05608) [gr-qc].

- [97] Ian Harry, Juan Calderón Bustillo, and Alex Nitz. “Searching for the Full Symphony of Black Hole Binary Mergers”. *Phys. Rev. D* 97.2 (2018), p. 023004. DOI: [10.1103/PhysRevD.97.023004](https://doi.org/10.1103/PhysRevD.97.023004). arXiv: [1709.09181](https://arxiv.org/abs/1709.09181) [gr-qc].
- [98] Juan Calderón Bustillo et al. “Impact of Gravitational Radiation Higher Order Modes on Single Aligned-Spin Gravitational Wave Searches for Binary Black Holes”. *Phys. Rev. D* 93.8 (2016), p. 084019. DOI: [10.1103/PhysRevD.93.084019](https://doi.org/10.1103/PhysRevD.93.084019). arXiv: [1511.02060](https://arxiv.org/abs/1511.02060) [gr-qc].
- [99] Tyson B. Littenberg et al. “Systematic Biases in Parameter Estimation of Binary Black-Hole Mergers”. *Physical Review D* 87.10 (2013), p. 104003. DOI: [10.1103/PhysRevD.87.104003](https://doi.org/10.1103/PhysRevD.87.104003). arXiv: [1210.0893](https://arxiv.org/abs/1210.0893) [gr-qc].
- [100] Feroz H. Shaik et al. “Impact of Subdominant Modes on the Interpretation of Gravitational-Wave Signals from Heavy Binary Black Hole Systems”. *Physical Review D* 101.12 (2020), p. 124054. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.101.124054](https://doi.org/10.1103/PhysRevD.101.124054). arXiv: [1911.02693](https://arxiv.org/abs/1911.02693) [gr-qc].
- [101] Chinmay Kalaghatgi, Mark Hannam, and Vivien Raymond. “Parameter Estimation with a Spinning Multimode Waveform Model”. *Physical Review D* 101.10 (2020), p. 103004. DOI: [10.1103/PhysRevD.101.103004](https://doi.org/10.1103/PhysRevD.101.103004). arXiv: [1909.10010](https://arxiv.org/abs/1909.10010) [gr-qc].
- [102] Marta Colleoni et al. “Towards the Routine Use of Subdominant Harmonics in Gravitational-Wave Inference: Re-Analysis of GW190412 with Generation X Waveform Models”. 103.2 (2021), p. 024029. DOI: [10.1103/PhysRevD.103.024029](https://doi.org/10.1103/PhysRevD.103.024029). arXiv: [2010.05830](https://arxiv.org/abs/2010.05830) [gr-qc].
- [103] K. S. Thorne. “GRAVITATIONAL RADIATION” (1987).
- [104] Patrick R. Brady and Stephen Fairhurst. “Interpreting the Results of Searches for Gravitational Waves from Coalescing Binaries”. *Classical and Quantum Gravity* 25.10 (2008), p. 105002. ISSN: 0264-9381. DOI: [10.1088/0264-9381/25/10/105002](https://doi.org/10.1088/0264-9381/25/10/105002). arXiv: [0707.2410](https://arxiv.org/abs/0707.2410) [gr-qc].

- [105] Lionel London et al. “First Higher-Multipole Model of Gravitational Waves from Spinning and Coalescing Black-Hole Binaries”. *Phys. Rev. Lett.* 120.16 (2018), p. 161102. DOI: [10.1103/PhysRevLett.120.161102](https://doi.org/10.1103/PhysRevLett.120.161102). arXiv: [1708.00404](https://arxiv.org/abs/1708.00404) [[gr-qc](#)].
- [106] Dragoljub Marković. “Possibility of Determining Cosmological Parameters from Measurements of Gravitational Waves Emitted by Coalescing, Compact Binaries”. *Physical Review D* 48.10 (1993), pp. 4738–4756. DOI: [10.1103/PhysRevD.48.4738](https://doi.org/10.1103/PhysRevD.48.4738).
- [107] Samaya Nissanke et al. “EXPLORING SHORT GAMMA-RAY BURSTS AS GRAVITATIONAL-WAVE STANDARD SIRENS”. *The Astrophysical Journal* 725.1 (2010), pp. 496–514. ISSN: 0004-637X. DOI: [10.1088/0004-637X/725/1/496](https://doi.org/10.1088/0004-637X/725/1/496). arXiv: [0904.1017](https://arxiv.org/abs/0904.1017) [[astro-ph.CO](#)].
- [108] Philip B. Graff, Alessandra Buonanno, and B. S. Sathyaprakash. “Missing Link: Bayesian Detection and Measurement of Intermediate-Mass Black-Hole Binaries”. *Phys. Rev. D* 92.2 (2015), p. 022002. DOI: [10.1103/PhysRevD.92.022002](https://doi.org/10.1103/PhysRevD.92.022002). arXiv: [1504.04766](https://arxiv.org/abs/1504.04766) [[gr-qc](#)].
- [109] R. Abbott, T.D. Abbott, S. Abraham, et al. “GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses”. *Physical Review D* 102.4 (2020), p. 043015. DOI: [10.1103/PhysRevD.102.043015](https://doi.org/10.1103/PhysRevD.102.043015). arXiv: [2004.08342](https://arxiv.org/abs/2004.08342) [[astro-ph.HE](#)].
- [110] Alejandro Bohé et al. “Improved Effective-One-Body Model of Spinning, Non-precessing Binary Black Holes for the Era of Gravitational-Wave Astrophysics with Advanced Detectors”. *Phys. Rev. D* 95.4 (2017), p. 044028. DOI: [10.1103/PhysRevD.95.044028](https://doi.org/10.1103/PhysRevD.95.044028). arXiv: [1611.03703](https://arxiv.org/abs/1611.03703) [[gr-qc](#)].
- [111] Roberto Cotesta et al. “Enriching the Symphony of Gravitational Waves from Binary Black Holes by Tuning Higher Harmonics”. *Phys. Rev. D* 98.8 (2018), p. 084028. DOI: [10.1103/PhysRevD.98.084028](https://doi.org/10.1103/PhysRevD.98.084028). arXiv: [1803.10701](https://arxiv.org/abs/1803.10701) [[gr-qc](#)].
- [112] Roberto Cotesta, Sylvain Marsat, and Michael Pürrer. “Frequency-Domain Reduced-Order Model of Aligned-Spin Effective-One-Body Waveforms with Higher-Order

- Modes”. *Physical Review D* 101.12 (2020), p. 124040. DOI: [10.1103/PhysRevD.101.124040](https://doi.org/10.1103/PhysRevD.101.124040). arXiv: [2003.12079 \[gr-qc\]](https://arxiv.org/abs/2003.12079).
- [113] Nathan Steinle and Michael Kesden. “Signatures of Spin Precession and Nutation in Isolated Black-Hole Binaries”. *Physical Review D* 106.6 (2022), p. 063028. DOI: [10.1103/PhysRevD.106.063028](https://doi.org/10.1103/PhysRevD.106.063028). arXiv: [2206.00391 \[astro-ph.HE\]](https://arxiv.org/abs/2206.00391).
- [114] Nathan K. Johnson-McDaniel, Sumeet Kulkarni, and Anuradha Gupta. “Inferring Spin Tilts at Formation from Gravitational Wave Observations of Binary Black Holes: Interfacing Precession-Averaged and Orbit-Averaged Spin Evolution”. *Physical Review D* 106.2 (2022), p. 023001. DOI: [10.1103/PhysRevD.106.023001](https://doi.org/10.1103/PhysRevD.106.023001). arXiv: [2107.11902 \[astro-ph.HE\]](https://arxiv.org/abs/2107.11902).
- [115] Juan Calderón Bustillo, Pablo Laguna, and Deirdre Shoemaker. “Detectability of Gravitational Waves from Binary Black Holes: Impact of Precession and Higher Modes”. *Phys. Rev. D* 95.10 (2017), p. 104038. DOI: [10.1103/PhysRevD.95.104038](https://doi.org/10.1103/PhysRevD.95.104038). arXiv: [1612.02340 \[gr-qc\]](https://arxiv.org/abs/1612.02340).
- [116] N. V. Krishnendu and Frank Ohme. “Interplay of Spin-Precession and Higher Harmonics in the Parameter Estimation of Binary Black Holes”. *Physical Review D* 105.6 (2022), p. 064012. DOI: [10.1103/PhysRevD.105.064012](https://doi.org/10.1103/PhysRevD.105.064012). arXiv: [2110.00766 \[gr-qc\]](https://arxiv.org/abs/2110.00766).
- [117] Sylvia Biscoveanu et al. “Measuring the Spins of Heavy Binary Black Holes”. *Physical Review D* 104.10 (2021), p. 103018. DOI: [10.1103/PhysRevD.104.103018](https://doi.org/10.1103/PhysRevD.104.103018). arXiv: [2106.06492 \[gr-qc\]](https://arxiv.org/abs/2106.06492).
- [118] Patricia Schmidt. “Studying and Modelling the Complete Gravitational-Wave Signal from Precessing Black Hole Binaries”. PhD thesis. Cardiff University, 2014.
- [119] Mark Hannam. “Modelling Gravitational Waves from Precessing Black-Hole Binaries: Progress, Challenges and Prospects”. *General Relativity and Gravitation* 46.9 (2014), p. 1767. ISSN: 1572-9532. DOI: [10.1007/s10714-014-1767-2](https://doi.org/10.1007/s10714-014-1767-2). arXiv: [1312.3641 \[gr-qc\]](https://arxiv.org/abs/1312.3641).

- [120] Antoni Ramos-Buades et al. “Validity of Common Modeling Approximations for Precessing Binary Black Holes with Higher-Order Modes”. *Phys. Rev. D* 101.10 (2020), p. 103014. DOI: [10.1103/PhysRevD.101.103014](https://doi.org/10.1103/PhysRevD.101.103014). arXiv: [2001.10936 \[gr-qc\]](https://arxiv.org/abs/2001.10936).
- [121] Katerina Chatziioannou et al. “Gravitational Waveforms for Precessing, Quasi-circular Compact Binaries with Multiple Scale Analysis: Small Spin Expansion”. *Physical Review D* 88.6 (2013), p. 063011. DOI: [10.1103/PhysRevD.88.063011](https://doi.org/10.1103/PhysRevD.88.063011). arXiv: [1307.4418 \[gr-qc\]](https://arxiv.org/abs/1307.4418).
- [122] Davide Gerosa et al. “Multi-Timescale Analysis of Phase Transitions in Precessing Black-Hole Binaries”. *Phys. Rev. D* 92 (2015), p. 064016. DOI: [10.1103/PhysRevD.92.064016](https://doi.org/10.1103/PhysRevD.92.064016). arXiv: [1506.03492 \[gr-qc\]](https://arxiv.org/abs/1506.03492).
- [123] Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. “Detecting Gravitational Waves from Precessing Binaries of Spinning Compact Objects: Adiabatic Limit”. *Phys. Rev. D* 67 (2003), p. 104025. DOI: [10.1103/PhysRevD.67.104025](https://doi.org/10.1103/PhysRevD.67.104025). arXiv: [gr-qc/0211087](https://arxiv.org/abs/gr-qc/0211087).
- [124] Chris Van Den Broeck et al. “Template Banks to Search for Compact Binaries with Spinning Components in Gravitational Wave Data”. *Physical Review D* 80.2 (2009), p. 024009. DOI: [10.1103/PhysRevD.80.024009](https://doi.org/10.1103/PhysRevD.80.024009). arXiv: [0904.1715 \[gr-qc\]](https://arxiv.org/abs/0904.1715).
- [125] Michael Boyle. “Angular Velocity of Gravitational Radiation from Precessing Binaries and the Corotating Frame”. *Phys. Rev. D* 87.10 (2013), p. 104006. DOI: [10.1103/PhysRevD.87.104006](https://doi.org/10.1103/PhysRevD.87.104006). arXiv: [1302.2919 \[gr-qc\]](https://arxiv.org/abs/1302.2919).
- [126] Stanislav Babak, Andrea Taracchini, and Alessandra Buonanno. “Validating the Effective-One-Body Model of Spinning, Precessing Binary Black Holes against Numerical Relativity”. *Phys. Rev. D* 95.2 (2017), p. 024010. DOI: [10.1103/PhysRevD.95.024010](https://doi.org/10.1103/PhysRevD.95.024010). arXiv: [1607.05661 \[gr-qc\]](https://arxiv.org/abs/1607.05661).
- [127] Alejandro Bohé et al. *PhenomPv2 – Technical Notes for the LAL Implementation*. 2016. URL: <https://dcc.ligo.org/LIGO-T1500602/public>.

- [128] Patricia Schmidt, Mark Hannam, and Sascha Husa. “Towards Models of Gravitational Waveforms from Generic Binaries: A Simple Approximate Mapping between Precessing and Non-Precessing Inspirational Signals”. *Phys. Rev. D* 86 (2012), p. 104063. DOI: [10.1103/PhysRevD.86.104063](https://doi.org/10.1103/PhysRevD.86.104063). arXiv: [1207.3088](https://arxiv.org/abs/1207.3088) [gr-qc].
- [129] Larne Pekowsky et al. “Comparing Gravitational Waves from Nonprecessing and Precessing Black Hole Binaries in the Corotating Frame”. *Phys. Rev. D* 88.2 (2013), p. 024040. DOI: [10.1103/PhysRevD.88.024040](https://doi.org/10.1103/PhysRevD.88.024040). arXiv: [1304.3176](https://arxiv.org/abs/1304.3176) [gr-qc].
- [130] Chinmay Kalaghatgi and Mark Hannam. “Investigating the effect of in-plane spin directions for precessing binary black hole systems”. *Phys. Rev. D* 103.2 (2021), p. 024024. DOI: [10.1103/PhysRevD.103.024024](https://doi.org/10.1103/PhysRevD.103.024024). arXiv: [2008.09957](https://arxiv.org/abs/2008.09957) [gr-qc].
- [131] Michael Boyle et al. “Gravitational-Wave Modes from Precessing Black-Hole Binaries” (2014). arXiv: [1409.4431](https://arxiv.org/abs/1409.4431) [gr-qc]. URL: <https://arxiv.org/abs/1409.4431>.
- [132] Davide Gerosa et al. “A Generalized Precession Parameter χ to Interpret Gravitational-Wave Data”. *Phys. Rev. D* 103.6 (2021), p. 064067. DOI: [10.1103/PhysRevD.103.064067](https://doi.org/10.1103/PhysRevD.103.064067). arXiv: [2011.11948](https://arxiv.org/abs/2011.11948) [gr-qc].
- [133] Sarp Akcay, Rossella Gamba, and Sebastiano Bernuzzi. “A Hybrid Post-Newtonian – Effective-One-Body Scheme for Spin-Precessing Compact-Binary Waveforms”. *Phys. Rev. D* 103.2 (2021), p. 024014. DOI: [10.1103/PhysRevD.103.024014](https://doi.org/10.1103/PhysRevD.103.024014). arXiv: [2005.05338](https://arxiv.org/abs/2005.05338) [gr-qc].
- [134] Benjamin 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 (2019), p. 031040. DOI: [10.1103/PhysRevX.9.031040](https://doi.org/10.1103/PhysRevX.9.031040). arXiv: [1811.12907](https://arxiv.org/abs/1811.12907) [astro-ph.HE].
- [135] Enrico Barausse, Etienne Racine, and Alessandra Buonanno. “Hamiltonian of a Spinning Test Particle in Curved Spacetime”. *Physical Review D* 80.10 (2009), p. 104025. DOI: [10.1103/PhysRevD.85.069904](https://doi.org/10.1103/PhysRevD.85.069904). arXiv: [0907.4745](https://arxiv.org/abs/0907.4745) [gr-qc].

- [136] Enrico Barausse and Alessandra Buonanno. “Improved Effective-One-Body Hamiltonian for Spinning Black-Hole Binaries”. *Physical Review D* 81.8 (2010), p. 084024. DOI: [10.1103/PhysRevD.81.084024](https://doi.org/10.1103/PhysRevD.81.084024). arXiv: [0912.3517](https://arxiv.org/abs/0912.3517) [gr-qc].
- [137] Thibault Damour and Alessandro Nagar. “Faithful Effective-One-Body Waveforms of Small-Mass-Ratio Coalescing Black Hole Binaries”. *Physical Review D* 76.6 (2007), p. 064028. DOI: [10.1103/PhysRevD.76.064028](https://doi.org/10.1103/PhysRevD.76.064028). arXiv: [0705.2519](https://arxiv.org/abs/0705.2519) [gr-qc].
- [138] Thibault Damour, Bala R. Iyer, and Alessandro Nagar. “Improved Resummation of Post-Newtonian Multipolar Waveforms from Circularized Compact Binaries”. *Physical Review D* 79.6 (2009), p. 064004. DOI: [10.1103/PhysRevD.79.064004](https://doi.org/10.1103/PhysRevD.79.064004). arXiv: [0811.2069](https://arxiv.org/abs/0811.2069) [gr-qc].
- [139] Yi Pan et al. “Post-Newtonian Factorized Multipolar Waveforms for Spinning, Non-Precessing Black-Hole Binaries”. *Physical Review D: Particles and Fields* 83 (2011), p. 064003. DOI: [10.1103/PhysRevD.83.064003](https://doi.org/10.1103/PhysRevD.83.064003). arXiv: [1006.0431](https://arxiv.org/abs/1006.0431) [gr-qc].
- [140] Piero Retegno et al. “Comparing Effective-One-Body Hamiltonians for Spin-Aligned Coalescing Binaries”. *Physical Review D* 101.10 (2020), p. 104027. DOI: [10.1103/PhysRevD.101.104027](https://doi.org/10.1103/PhysRevD.101.104027). arXiv: [1911.10818](https://arxiv.org/abs/1911.10818) [gr-qc].
- [141] Mohammed Khalil et al. “Fourth Post-Newtonian Effective-One-Body Hamiltonians with Generic Spins”. *Physical Review D* 101.10 (2020), p. 104034. DOI: [10.1103/PhysRevD.101.104034](https://doi.org/10.1103/PhysRevD.101.104034). arXiv: [2003.04469](https://arxiv.org/abs/2003.04469) [gr-qc].
- [142] Thibault Damour, Bala R. Iyer, and B. S. Sathyaprakash. “Comparison of Search Templates for Gravitational Waves from Binary Inspiral: 3.5PN Update”. *Physical Review D* 66.2 (2002), p. 027502. DOI: [10.1103/PhysRevD.66.027502](https://doi.org/10.1103/PhysRevD.66.027502). arXiv: [gr-qc/0207021](https://arxiv.org/abs/gr-qc/0207021).
- [143] K. G. Arun et al. “Parameter Estimation of Inspiralling Compact Binaries Using 3.5 Post-Newtonian Gravitational Wave Phasing: The Nonspinning Case”. *Physical Review D* 71.8 (2005), p. 084008. DOI: [10.1103/PhysRevD.71.084008](https://doi.org/10.1103/PhysRevD.71.084008). arXiv: [gr-qc/0411146](https://arxiv.org/abs/gr-qc/0411146).

- [144] Enno Harms et al. “A New Gravitational Wave Generation Algorithm for Particle Perturbations of the Kerr Spacetime”. *Classical and Quantum Gravity* 31.24 (2014), p. 245004. ISSN: 0264-9381. DOI: [10.1088/0264-9381/31/24/245004](https://doi.org/10.1088/0264-9381/31/24/245004). arXiv: [1406.5983](https://arxiv.org/abs/1406.5983) [gr-qc].
- [145] Xisco Jiménez-Forteza et al. “Hierarchical Data-Driven Approach to Fitting Numerical Relativity Data for Nonprecessing Binary Black Holes with an Application to Final Spin and Radiated Energy”. *Phys. Rev. D* 95.6 (2017), p. 064024. DOI: [10.1103/PhysRevD.95.064024](https://doi.org/10.1103/PhysRevD.95.064024). arXiv: [1611.00332](https://arxiv.org/abs/1611.00332) [gr-qc].
- [146] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. “Tail-Induced Spin-Orbit Effect in the Gravitational Radiation of Compact Binaries”. *Physical Review D* 84.6 (2011), p. 064041. DOI: [10.1103/PhysRevD.84.064041](https://doi.org/10.1103/PhysRevD.84.064041). arXiv: [1104.5659](https://arxiv.org/abs/1104.5659) [gr-qc].
- [147] Sylvain Marsat et al. “Gravitational Waves from Spinning Compact Object Binaries: New Post-Newtonian Results” (2013). DOI: [10.48550/arXiv.1312.5375](https://doi.org/10.48550/arXiv.1312.5375). arXiv: [1312.5375](https://arxiv.org/abs/1312.5375) [gr-qc].
- [148] Katerina Chatziioannou et al. “Constructing Gravitational Waves from Generic Spin-Precessing Compact Binary Inspirals”. *Physical Review D* 95.10 (2017), p. 104004. DOI: [10.1103/PhysRevD.95.104004](https://doi.org/10.1103/PhysRevD.95.104004). arXiv: [1703.03967](https://arxiv.org/abs/1703.03967) [gr-qc].
- [149] Antoine Klein, Neil Cornish, and Nicolás Yunes. “Gravitational Waveforms for Precessing, Quasicircular Binaries via Multiple Scale Analysis and Uniform Asymptotics: The near Spin Alignment Case”. *Physical Review D* 88.12 (2013), p. 124015. DOI: [10.1103/PhysRevD.88.124015](https://doi.org/10.1103/PhysRevD.88.124015). arXiv: [1305.1932](https://arxiv.org/abs/1305.1932) [gr-qc].
- [150] Michael Kesden et al. “Effective Potentials and Morphological Transitions for Binary Black Hole Spin Precession”. *Physical Review Letters* 114.8 (2015), p. 081103. DOI: [10.1103/PhysRevLett.114.081103](https://doi.org/10.1103/PhysRevLett.114.081103). arXiv: [1411.0674](https://arxiv.org/abs/1411.0674) [gr-qc].

- [151] Scott E. Field et al. “Reduced Basis Catalogs for Gravitational Wave Templates”. *Phys. Rev. Lett.* 106 (2011), p. 221102. DOI: [10.1103/PhysRevLett.106.221102](https://doi.org/10.1103/PhysRevLett.106.221102). arXiv: [1101.3765 \[gr-qc\]](https://arxiv.org/abs/1101.3765).
- [152] Scott E. Field et al. “Fast Prediction and Evaluation of Gravitational Waveforms Using Surrogate Models”. *Phys. Rev. X* 4.3 (2014), p. 031006. DOI: [10.1103/PhysRevX.4.031006](https://doi.org/10.1103/PhysRevX.4.031006). arXiv: [1308.3565 \[gr-qc\]](https://arxiv.org/abs/1308.3565).
- [153] Sebastian Khan and Rhys Green. “Gravitational-Wave Surrogate Models Powered by Artificial Neural Networks”. *Phys. Rev. D* 103.6 (2021), p. 064015. DOI: [10.1103/PhysRevD.103.064015](https://doi.org/10.1103/PhysRevD.103.064015). arXiv: [2008.12932 \[gr-qc\]](https://arxiv.org/abs/2008.12932).
- [154] Styliani-Christina Fragkouli et al. “Deep Residual Error and Bag-of-Tricks Learning for Gravitational Wave Surrogate Modeling” (2022). arXiv: [2203.08434 \[astro-ph.IM\]](https://arxiv.org/abs/2203.08434). URL: <https://arxiv.org/abs/2203.08434>.
- [155] Alessandro Nagar et al. “Time-Domain Effective-One-Body Gravitational Waveforms for Coalescing Compact Binaries with Nonprecessing Spins, Tides and Self-Spin Effects”. *Phys. Rev. D* 98.10 (2018), p. 104052. DOI: [10.1103/PhysRevD.98.104052](https://doi.org/10.1103/PhysRevD.98.104052). arXiv: [1806.01772 \[gr-qc\]](https://arxiv.org/abs/1806.01772).
- [156] Stefano Schmidt et al. “Machine Learning Gravitational Waves from Binary Black Hole Mergers”. *Phys. Rev. D* 103.4 (2021), p. 043020. DOI: [10.1103/PhysRevD.103.043020](https://doi.org/10.1103/PhysRevD.103.043020). arXiv: [2011.01958 \[gr-qc\]](https://arxiv.org/abs/2011.01958).
- [157] Chad R. Galley and Patricia Schmidt. “Fast and Efficient Evaluation of Gravitational Waveforms via Reduced-Order Spline Interpolation” (2016). arXiv: [1611.07529 \[gr-qc\]](https://arxiv.org/abs/1611.07529). URL: <https://arxiv.org/abs/1611.07529>.
- [158] Maxime Barrault et al. “An ‘Empirical Interpolation’ Method: Application to Efficient Reduced-Basis Discretization of Partial Differential Equations”. *Comptes Rendus Mathematique* 339.9 (2004), pp. 667–672. ISSN: 1631-073X. DOI: [10.1016/j.crma.2004.08.006](https://doi.org/10.1016/j.crma.2004.08.006).

-
- [159] Yvon Maday et al. “A General Multipurpose Interpolation Procedure: The Magic Points”. *Communications on Pure and Applied Analysis* 8.1 (2008), pp. 383–404. DOI: [10.3934/cpaa.2009.8.383](https://doi.org/10.3934/cpaa.2009.8.383).
- [160] Vijay Varma et al. “Surrogate Models for Precessing Binary Black Hole Simulations with Unequal Masses”. *Phys. Rev. Research*. 1 (2019), p. 033015. DOI: [10.1103/PhysRevResearch.1.033015](https://doi.org/10.1103/PhysRevResearch.1.033015). arXiv: [1905.09300 \[gr-qc\]](https://arxiv.org/abs/1905.09300).
- [161] Vijay Varma et al. “Surrogate Model of Hybridized Numerical Relativity Binary Black Hole Waveforms”. *Phys. Rev. D* 99.6 (2019), p. 064045. DOI: [10.1103/PhysRevD.99.064045](https://doi.org/10.1103/PhysRevD.99.064045). arXiv: [1812.07865 \[gr-qc\]](https://arxiv.org/abs/1812.07865).
- [162] Daniel Williams et al. “Precessing Numerical Relativity Waveform Surrogate Model for Binary Black Holes: A Gaussian Process Regression Approach”. *Phys. Rev. D* 101.6 (2020), p. 063011. DOI: [10.1103/PhysRevD.101.063011](https://doi.org/10.1103/PhysRevD.101.063011). arXiv: [1903.09204 \[gr-qc\]](https://arxiv.org/abs/1903.09204).
- [163] Tousif Islam et al. “Eccentric Binary Black Hole Surrogate Models for the Gravitational Waveform and Remnant Properties: Comparable Mass, Nonspinning Case”. *Phys. Rev. D* 103.6 (2021), p. 064022. DOI: [10.1103/PhysRevD.103.064022](https://doi.org/10.1103/PhysRevD.103.064022). arXiv: [2101.11798 \[gr-qc\]](https://arxiv.org/abs/2101.11798).
- [164] Jooheon Yoo et al. “Targeted Large Mass Ratio Numerical Relativity Surrogate Waveform Model for GW190814”. 106.4 (2022), p. 044001. DOI: [10.1103/PhysRevD.106.044001](https://doi.org/10.1103/PhysRevD.106.044001). arXiv: [2203.10109 \[gr-qc\]](https://arxiv.org/abs/2203.10109).
- [165] Mark Hannam et al. “Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms”. *Phys. Rev. Lett.* 113.15 (2014), p. 151101. DOI: [10.1103/PhysRevLett.113.151101](https://doi.org/10.1103/PhysRevLett.113.151101). arXiv: [1308.3271 \[gr-qc\]](https://arxiv.org/abs/1308.3271).
- [166] Yi Pan et al. “Inspirals-Merger-Ringdown Waveforms of Spinning, Precessing Black-Hole Binaries in the Effective-One-Body Formalism”. *Physical Review D* 89.8 (2014), p. 084006. DOI: [10.1103/PhysRevD.89.084006](https://doi.org/10.1103/PhysRevD.89.084006). arXiv: [1307.6232 \[gr-qc\]](https://arxiv.org/abs/1307.6232).

- [167] Sebastian Khan et al. “Including Higher Order Multipoles in Gravitational-Wave Models for Precessing Binary Black Holes”. *Phys. Rev. D* 101.2 (2020), p. 024056. DOI: [10.1103/PhysRevD.101.024056](https://doi.org/10.1103/PhysRevD.101.024056). arXiv: [1911.06050](https://arxiv.org/abs/1911.06050) [gr-qc].
- [168] Geraint Pratten et al. “Computationally Efficient Models for the Dominant and Subdominant Harmonic Modes of Precessing Binary Black Holes”. *Phys. Rev. D* 103.10 (2021), p. 104056. DOI: [10.1103/PhysRevD.103.104056](https://doi.org/10.1103/PhysRevD.103.104056). arXiv: [2004.06503](https://arxiv.org/abs/2004.06503) [gr-qc].
- [169] David Reitze et al. “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO”. 51.7 (2019), p. 035. arXiv: [1907.04833](https://arxiv.org/abs/1907.04833) [astro-ph.IM]. URL: <https://arxiv.org/abs/1907.04833>.
- [170] Michael Pürrer and Carl-Johan Haster. “Gravitational Waveform Accuracy Requirements for Future Ground-Based Detectors”. *Phys. Rev. Res.* 2.2 (2020), p. 023151. DOI: [10.1103/PhysRevResearch.2.023151](https://doi.org/10.1103/PhysRevResearch.2.023151). arXiv: [1912.10055](https://arxiv.org/abs/1912.10055) [gr-qc].
- [171] Abdul H. Mroue et al. “Catalog of 174 Binary Black Hole Simulations for Gravitational Wave Astronomy”. *Phys. Rev. Lett.* 111.24 (2013), p. 241104. DOI: [10.1103/PhysRevLett.111.241104](https://doi.org/10.1103/PhysRevLett.111.241104). arXiv: [1304.6077](https://arxiv.org/abs/1304.6077) [gr-qc].
- [172] Leonardo "Werneck and others". *The Einstein Toolkit*. Version The "Karl Schwarzschild" release, ET_2023_05. 2023. DOI: [10.5281/zenodo.7942541](https://doi.org/10.5281/zenodo.7942541). URL: <https://doi.org/10.5281/zenodo.7942541>.
- [173] Patricia Schmidt et al. “Tracking the Precession of Compact Binaries from Their Gravitational-Wave Signal”. *Phys. Rev. D* 84 (2011), p. 024046. DOI: [10.1103/PhysRevD.84.024046](https://doi.org/10.1103/PhysRevD.84.024046). arXiv: [1012.2879](https://arxiv.org/abs/1012.2879) [gr-qc].
- [174] Patricia Schmidt, Frank Ohme, and Mark Hannam. “Towards Models of Gravitational Waveforms from Generic Binaries II: Modelling Precession Effects with a Single Effective Precession Parameter”. *Phys. Rev. D* 91.2 (2015), p. 024043. DOI: [10.1103/PhysRevD.91.024043](https://doi.org/10.1103/PhysRevD.91.024043). arXiv: [1408.1810](https://arxiv.org/abs/1408.1810) [gr-qc].

- [175] Jonathan Blackman et al. *Gwsurrogate*. URL: <https://pypi.python.org/pypi/gwsurrogate/>.
- [176] Collin Capano, Yi Pan, and Alessandra Buonanno. “Impact of Higher Harmonics in Searching for Gravitational Waves from Nonspinning Binary Black Holes”. *Phys. Rev. D* 89.10 (2014), p. 102003. DOI: [10.1103/PhysRevD.89.102003](https://doi.org/10.1103/PhysRevD.89.102003). arXiv: [1311.1286 \[gr-qc\]](https://arxiv.org/abs/1311.1286).
- [177] William H. Press. “Long Wave Trains of Gravitational Waves from a Vibrating Black Hole”. *Astrophys. J. Lett.* 170 (1971), pp. L105–L108. DOI: [10.1086/180849](https://doi.org/10.1086/180849).
- [178] S. Chandrasekhar and Steven L. Detweiler. “The Quasi-Normal Modes of the Schwarzschild Black Hole”. *Proc. Roy. Soc. Lond. A* 344 (1975), pp. 441–452. DOI: [10.1098/rspa.1975.0112](https://doi.org/10.1098/rspa.1975.0112).
- [179] Steven L. Detweiler. “Black Holes and Gravitational Waves. III. The Resonant Frequencies of Rotating Holes”. *Astrophys. J.* 239 (1980), pp. 292–295. DOI: [10.1086/158109](https://doi.org/10.1086/158109).
- [180] Kostas D. Kokkotas and Bernd G. Schmidt. “Quasinormal Modes of Stars and Black Holes”. *Living Rev. Rel.* 2 (1999), p. 2. DOI: [10.12942/lrr-1999-2](https://doi.org/10.12942/lrr-1999-2). arXiv: [gr-qc/9909058](https://arxiv.org/abs/gr-qc/9909058).
- [181] Vijay Varma et al. “High-Accuracy Mass, Spin, and Recoil Predictions of Generic Black-Hole Merger Remnants”. *Phys. Rev. Lett.* 122.1 (2019), p. 011101. DOI: [10.1103/PhysRevLett.122.011101](https://doi.org/10.1103/PhysRevLett.122.011101). arXiv: [1809.09125 \[gr-qc\]](https://arxiv.org/abs/1809.09125).
- [182] Vijay Varma and Leo C. Stein. *Vijayvarma392/surfinBH: Surrogate Final BH Properties*. Zenodo. 2018. DOI: [10.5281/zenodo.1418525](https://doi.org/10.5281/zenodo.1418525).
- [183] Fabian Hofmann, Enrico Barausse, and Luciano Rezzolla. “The Final Spin from Binary Black Holes in Quasi-Circular Orbits”. *Astrophys. J. Lett.* 825.2 (2016), p. L19. DOI: [10.3847/2041-8205/825/2/L19](https://doi.org/10.3847/2041-8205/825/2/L19). arXiv: [1605.01938 \[gr-qc\]](https://arxiv.org/abs/1605.01938).

- [184] Alessandra Buonanno et al. “A Quasi-physical Family of Gravity-Wave Templates for Precessing Binaries of Spinning Compact Objects. 2. Application to Double-Spin Precessing Binaries”. *Phys. Rev. D* 70 (2004), p. 104003. DOI: [10.1103/PhysRevD.74.029902](https://doi.org/10.1103/PhysRevD.74.029902). arXiv: [gr-qc/0405090](https://arxiv.org/abs/gr-qc/0405090).
- [185] Jose A. Gonzalez et al. “Total Recoil: The Maximum Kick from Nonspinning Black-Hole Binary Inspiral”. *Phys. Rev. Lett.* 98 (2007), p. 091101. DOI: [10.1103/PhysRevLett.98.091101](https://doi.org/10.1103/PhysRevLett.98.091101). arXiv: [gr-qc/0610154](https://arxiv.org/abs/gr-qc/0610154).
- [186] Bernd Brügmann et al. “Exploring Black Hole Superkicks”. *Phys. Rev. D* 77 (2008), p. 124047. DOI: [10.1103/PhysRevD.77.124047](https://doi.org/10.1103/PhysRevD.77.124047). arXiv: [0707.0135 \[gr-qc\]](https://arxiv.org/abs/0707.0135).
- [187] Manuela Campanelli et al. “Maximum Gravitational Recoil”. *Physical Review Letters* 98 (2007), p. 231102. DOI: [10.1103/PhysRevLett.98.231102](https://doi.org/10.1103/PhysRevLett.98.231102). arXiv: [gr-qc/0702133](https://arxiv.org/abs/gr-qc/0702133).
- [188] John G. Baker et al. “Modeling Kicks from the Merger of Generic Black-Hole Binaries”. *Astrophys. J. Lett.* 682 (2008), pp. L29–L32. DOI: [10.1086/590927](https://doi.org/10.1086/590927). arXiv: [0802.0416 \[astro-ph\]](https://arxiv.org/abs/0802.0416).
- [189] Galley, Chad R. *RomPy*. URL: <https://bitbucket.org/chadgalley/rompy/src/master/>.
- [190] J J Hopfield. “Neural Networks and Physical Systems with Emergent Collective Computational Abilities.” *Proceedings of the National Academy of Sciences* 79.8 (1982), pp. 2554–2558. DOI: [10.1073/pnas.79.8.2554](https://doi.org/10.1073/pnas.79.8.2554).
- [191] Charu C. Aggarwal. *Neural Networks and Deep Learning: A Textbook*. Cham: Springer International Publishing, 2018. ISBN: 978-3-319-94463-0. DOI: [10.1007/978-3-319-94463-0](https://doi.org/10.1007/978-3-319-94463-0).
- [192] Ian Harry et al. “Searching for Gravitational Waves from Compact Binaries with Precessing Spins”. *Phys. Rev. D* 94.2 (2016), p. 024012. DOI: [10.1103/PhysRevD.94.024012](https://doi.org/10.1103/PhysRevD.94.024012). arXiv: [1603.02444 \[gr-qc\]](https://arxiv.org/abs/1603.02444).

- [193] F. Pedregosa et al. “Scikit-Learn: Machine Learning in Python”. 12 (2011), pp. 2825–2830.
- [194] Martín Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. 2015. URL: <https://www.tensorflow.org/>.
- [195] François Chollet et al. *Keras*. 2015. URL: <https://keras.io>.
- [196] James Bergstra, Daniel Yamins, and David Cox. “Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures”. *International Conference on Machine Learning*. PMLR, 2013, pp. 115–123.
- [197] Kunihiko Fukushima. “Visual Feature Extraction by a Multilayered Network of Analog Threshold Elements”. 5.4 (1969), pp. 322–333.
- [198] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. “Deep Sparse Rectifier Neural Networks”. *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. JMLR Workshop and Conference Proceedings, 2011, pp. 315–323.
- [199] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. *Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs)*. 2015. DOI: [10.48550/ARXIV.1511.07289](https://doi.org/10.48550/ARXIV.1511.07289).
- [200] David L. Elliott. *A Better Activation Function for Artificial Neural Networks*. 1993.
- [201] Diederik P Kingma and Jimmy Ba. “Adam: A Method for Stochastic Optimization” (2014). arXiv: [1412.6980](https://arxiv.org/abs/1412.6980).
- [202] Timothy Dozat. “Incorporating Nesterov Momentum into Adam”. 1:2013-2016 (2016).
- [203] Matthew D. Zeiler. “ADADELTA: An Adaptive Learning Rate Method” (2012). DOI: [10.48550/ARXIV.1212.5701](https://doi.org/10.48550/ARXIV.1212.5701).

- [204] Johannes Lederer. *Activation Functions in Artificial Neural Networks: A Systematic Overview*. 2021. DOI: [10.48550/ARXIV.2101.09957](https://doi.org/10.48550/ARXIV.2101.09957).
- [205] Sebastian Ruder. *An Overview of Gradient Descent Optimization Algorithms*. 2016. DOI: [10.48550/ARXIV.1609.04747](https://doi.org/10.48550/ARXIV.1609.04747). arXiv: [1609.04747](https://arxiv.org/abs/1609.04747) [[cs.LG](#)].
- [206] Michael Boyle, Robert Owen, and Harald P. Pfeiffer. “A Geometric Approach to the Precession of Compact Binaries”. *Phys. Rev. D* 84 (2011), p. 124011. DOI: [10.1103/PhysRevD.84.124011](https://doi.org/10.1103/PhysRevD.84.124011). arXiv: [1110.2965](https://arxiv.org/abs/1110.2965) [[gr-qc](#)].
- [207] Benjamin P. Abbott et al. “Prospects for Observing and Localizing Gravitational-Wave Transients with Advanced LIGO, Advanced Virgo and KAGRA”. *Living Rev. Rel.* 23.1 (2020), p. 3. DOI: [10.1007/s41114-020-00026-9](https://doi.org/10.1007/s41114-020-00026-9). arXiv: [1304.0670](https://arxiv.org/abs/1304.0670) [[gr-qc](#)].
- [208] Bhooshan Gadre et al. “A Fully Precessing Higher-Mode Surrogate Model of Effective-One-Body Waveforms” (2022). arXiv: [2203.00381](https://arxiv.org/abs/2203.00381) [[gr-qc](#)]. URL: <https://arxiv.org/abs/2203.00381>.
- [209] Alvin J. K. Chua and Michele Vallisneri. “Learning Bayesian Posteriors with Neural Networks for Gravitational-Wave Inference”. *Phys. Rev. Lett.* 124.4 (2020), p. 041102. DOI: [10.1103/PhysRevLett.124.041102](https://doi.org/10.1103/PhysRevLett.124.041102). arXiv: [1909.05966](https://arxiv.org/abs/1909.05966) [[gr-qc](#)].
- [210] Richard Abbott et al. “Properties and Astrophysical Implications of the 150 M \odot Binary Black Hole Merger GW190521”. *Astrophys. J. Lett.* 900.1 (2020), p. L13. DOI: [10.3847/2041-8213/aba493](https://doi.org/10.3847/2041-8213/aba493). arXiv: [2009.01190](https://arxiv.org/abs/2009.01190) [[astro-ph.HE](#)].
- [211] Isobel M. Romero-Shaw et al. “GW190521: orbital eccentricity and signatures of dynamical formation in a binary black hole merger signal”. *Astrophys. J. Lett.* 903.1 (2020), p. L5. DOI: [10.3847/2041-8213/abbe26](https://doi.org/10.3847/2041-8213/abbe26). arXiv: [2009.04771](https://arxiv.org/abs/2009.04771) [[astro-ph.HE](#)].
- [212] Juan Calderón Bustillo et al. “GW190521 as a Merger of Proca Stars: A Potential New Vector Boson of 8.7×10^{-13} eV”. *Phys. Rev. Lett.* 126.8 (2021), p. 081101. DOI: [10.1103/PhysRevLett.126.081101](https://doi.org/10.1103/PhysRevLett.126.081101). arXiv: [2009.05376](https://arxiv.org/abs/2009.05376) [[gr-qc](#)].

-
- [213] Alexander H. Nitz and Collin D. Capano. “GW190521 may be an intermediate mass ratio inspiral”. *Astrophys. J. Lett.* 907.1 (2021), p. L9. DOI: [10.3847/2041-8213/abccc5](https://doi.org/10.3847/2041-8213/abccc5). arXiv: [2010.12558](https://arxiv.org/abs/2010.12558) [[astro-ph.HE](#)].
- [214] Jeremy Sakstein et al. “Beyond the Standard Model Explanations of GW190521”. *Phys. Rev. Lett.* 125.26 (2020), p. 261105. DOI: [10.1103/PhysRevLett.125.261105](https://doi.org/10.1103/PhysRevLett.125.261105). arXiv: [2009.01213](https://arxiv.org/abs/2009.01213) [[gr-qc](#)].
- [215] Juan Calderón Bustillo et al. “Confusing Head-On Collisions with Precessing Intermediate-Mass Binary Black Hole Mergers”. *Phys. Rev. Lett.* 126.20 (2021), p. 201101. DOI: [10.1103/PhysRevLett.126.201101](https://doi.org/10.1103/PhysRevLett.126.201101). arXiv: [2009.01066](https://arxiv.org/abs/2009.01066) [[gr-qc](#)].
- [216] Rossella Gamba et al. “GW190521 as a dynamical capture of two nonspinning black holes”. *Nature Astron.* 7.1 (2023), pp. 11–17. DOI: [10.1038/s41550-022-01813-w](https://doi.org/10.1038/s41550-022-01813-w). arXiv: [2106.05575](https://arxiv.org/abs/2106.05575) [[gr-qc](#)].
- [217] Ethan Payne et al. “Curious case of GW200129: Interplay between spin-precession inference and data-quality issues”. *Phys. Rev. D* 106.10 (2022), p. 104017. DOI: [10.1103/PhysRevD.106.104017](https://doi.org/10.1103/PhysRevD.106.104017). arXiv: [2206.11932](https://arxiv.org/abs/2206.11932) [[gr-qc](#)].
- [218] Richard Abbott et al. “GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run” (2021). arXiv: [2111.03606](https://arxiv.org/abs/2111.03606) [[gr-qc](#)]. URL: <https://arxiv.org/abs/2111.03606>.
- [219] Michele Maggiore et al. “Science Case for the Einstein Telescope”. *JCAP* 03 (2020), p. 050. DOI: [10.1088/1475-7516/2020/03/050](https://doi.org/10.1088/1475-7516/2020/03/050). arXiv: [1912.02622](https://arxiv.org/abs/1912.02622) [[astro-ph.CO](#)].
- [220] B. S. Sathyaprakash et al. “Cosmology and the Early Universe” (2019). arXiv: [1903.09260](https://arxiv.org/abs/1903.09260) [[astro-ph.HE](#)].
- [221] Jonathan R. Gair et al. “Exploring intermediate and massive black-hole binaries with the Einstein Telescope”. *Gen. Rel. Grav.* 43 (2011), pp. 485–518. DOI: [10.1007/s10714-010-1104-3](https://doi.org/10.1007/s10714-010-1104-3). arXiv: [0907.5450](https://arxiv.org/abs/0907.5450) [[astro-ph.CO](#)].

- [222] Pau Amaro-Seoane and Lucia Santamaria. “Detection of IMBHs with ground-based gravitational wave observatories: A biography of a binary of black holes, from birth to death”. *Astrophys. J.* 722 (2010), pp. 1197–1206. DOI: [10.1088/0004-637X/722/2/1197](https://doi.org/10.1088/0004-637X/722/2/1197). arXiv: [0910.0254](https://arxiv.org/abs/0910.0254) [[astro-ph.CO](#)].
- [223] Davide Gerosa and Maya Fishbach. “Hierarchical Mergers of Stellar-Mass Black Holes and their Gravitational-Wave Signatures”. *Nature Astron.* 5.8 (2021), pp. 749–760. DOI: [10.1038/s41550-021-01398-w](https://doi.org/10.1038/s41550-021-01398-w). arXiv: [2105.03439](https://arxiv.org/abs/2105.03439) [[astro-ph.HE](#)].
- [224] S. Hild et al. “Sensitivity Studies for Third-Generation Gravitational Wave Observatories”. *Classical and Quantum Gravity* 28.9 (2011), p. 094013. ISSN: 0264-9381. DOI: [10.1088/0264-9381/28/9/094013](https://doi.org/10.1088/0264-9381/28/9/094013). arXiv: [1012.0908](https://arxiv.org/abs/1012.0908) [[gr-qc](#)].
- [225] Benjamin P. Abbott et al. “Exploring the Sensitivity of Next Generation Gravitational Wave Detectors”. *Class. Quant. Grav.* 34.4 (2017), p. 044001. DOI: [10.1088/1361-6382/aa51f4](https://doi.org/10.1088/1361-6382/aa51f4). arXiv: [1607.08697](https://arxiv.org/abs/1607.08697) [[astro-ph.IM](#)].
- [226] M. Punturo et al. “The Einstein Telescope: A third-generation gravitational wave observatory”. *Class. Quant. Grav.* 27 (2010). Ed. by Fulvio Ricci, p. 194002. DOI: [10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002).
- [227] Reed Essick, Salvatore Vitale, and Matthew Evans. “Frequency-dependent responses in third generation gravitational-wave detectors”. *Phys. Rev. D* 96.8 (2017), p. 084004. DOI: [10.1103/PhysRevD.96.084004](https://doi.org/10.1103/PhysRevD.96.084004). arXiv: [1708.06843](https://arxiv.org/abs/1708.06843) [[gr-qc](#)].
- [228] Irving John Good. *The Estimation of Probabilities: an Essay on Modern Bayesian methods*. MIT Press, 1965. ISBN: 978-0-2625-7015-2.
- [229] D. V. Lindley and A. F. M. Smith. “Bayes Estimates for the Linear Model”. *Journal of the Royal Statistical Society. Series B (Methodological)* 34.1 (1972), pp. 1–41. ISSN: 00359246. URL: <http://www.jstor.org/stable/2985048> (visited on 06/16/2023).
- [230] J. M. Bernardo et al. *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting*. Oxford University Press, 1992. ISBN: 978-0-19-852266-9.

-
- [231] David MacKay. “Comparison of Approximate Methods for Handling Hyperparameters”. *Neural Computation* 11.5 (1999). DOI: [10.1162/089976699300016331](https://doi.org/10.1162/089976699300016331).
- [232] Matthew R. Adams, Neil J. Cornish, and Tyson B. Littenberg. “Astrophysical Model Selection in Gravitational Wave Astronomy”. *Phys. Rev. D* 86 (2012), p. 124032. DOI: [10.1103/PhysRevD.86.124032](https://doi.org/10.1103/PhysRevD.86.124032). arXiv: [1209.6286](https://arxiv.org/abs/1209.6286) [[gr-qc](#)].
- [233] Ilya Mandel, Will M. Farr, and Jonathan R. Gair. “Extracting Distribution Parameters from Multiple Uncertain Observations with Selection Biases”. *Mon. Not. Roy. Astron. Soc.* 486.1 (2019), pp. 1086–1093. DOI: [10.1093/mnras/stz896](https://doi.org/10.1093/mnras/stz896). arXiv: [1809.02063](https://arxiv.org/abs/1809.02063) [[physics.data-an](#)].
- [234] Sebastian M. Gaebel et al. “Digging the Population of Compact Binary Mergers Out of the Noise”. *Mon. Not. Roy. Astron. Soc.* 484.3 (2019), pp. 4008–4023. DOI: [10.1093/mnras/stz225](https://doi.org/10.1093/mnras/stz225). arXiv: [1809.03815](https://arxiv.org/abs/1809.03815) [[astro-ph.IM](#)].
- [235] Salvatore Vitale et al. “Inferring the Properties of a Population of Compact Binaries in Presence of Selection Effects” (2020). DOI: [10.1007/978-981-15-4702-7_45-1](https://doi.org/10.1007/978-981-15-4702-7_45-1). arXiv: [2007.05579](https://arxiv.org/abs/2007.05579) [[astro-ph.IM](#)].
- [236] Simona Miller, Thomas A. Callister, and Will Farr. “The Low Effective Spin of Binary Black Holes and Implications for Individual Gravitational-Wave Events”. *Astrophys. J.* 895.2 (2020), p. 128. DOI: [10.3847/1538-4357/ab80c0](https://doi.org/10.3847/1538-4357/ab80c0). arXiv: [2001.06051](https://arxiv.org/abs/2001.06051) [[astro-ph.HE](#)].
- [237] Alejandro Bohe et al. “Next-to-next-to-leading order spin-orbit effects in the near-zone metric and precession equations of compact binaries”. *Class. Quant. Grav.* 30 (2013), p. 075017. DOI: [10.1088/0264-9381/30/7/075017](https://doi.org/10.1088/0264-9381/30/7/075017). arXiv: [1212.5520](https://arxiv.org/abs/1212.5520) [[gr-qc](#)].
- [238] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. “Higher-order spin effects in the dynamics of compact binaries. II. Radiation field”. *Phys. Rev. D* 74 (2006). [Erratum: *Phys.Rev.D* 75, 049903 (2007), Erratum: *Phys.Rev.D* 81, 089901 (2010)], p. 104034. DOI: [10.1103/PhysRevD.81.089901](https://doi.org/10.1103/PhysRevD.81.089901). arXiv: [gr-qc/0605140](https://arxiv.org/abs/gr-qc/0605140).

- [239] James M. Bardeen, William H. Press, and Saul A. Teukolsky. “Rotating Black Holes: Locally Nonrotating Frames, Energy Extraction, and Scalar Synchrotron Radiation”. 178 (1972), p. 347. DOI: [10.1086/151796](https://doi.org/10.1086/151796).
- [240] Lorenzo Speri et al. “A Roadmap of Gravitational Wave Data Analysis”. *Nature Astron.* 6.12 (2022), pp. 1356–1363. DOI: [10.1038/s41550-022-01849-y](https://doi.org/10.1038/s41550-022-01849-y).
- [241] Michele Vallisneri. “Use and Abuse of the Fisher Information Matrix in the Assessment of Gravitational-Wave Parameter-Estimation Prospects”. *Phys. Rev. D* 77 (2008), p. 042001. DOI: [10.1103/PhysRevD.77.042001](https://doi.org/10.1103/PhysRevD.77.042001). arXiv: [gr-qc/0703086](https://arxiv.org/abs/gr-qc/0703086).
- [242] Lee Samuel Finn. “Aperture Synthesis for Gravitational Wave Data Analysis: Deterministic Sources”. *Phys. Rev. D* 63 (2001), p. 102001. DOI: [10.1103/PhysRevD.63.102001](https://doi.org/10.1103/PhysRevD.63.102001). arXiv: [gr-qc/0010033](https://arxiv.org/abs/gr-qc/0010033).
- [243] Fausto Acernese et al. “Advanced Virgo: A Second-Generation Interferometric Gravitational Wave Detector”. *Class. Quant. Grav.* 32.2 (2014), p. 024001. DOI: [10.1088/0264-9381/32/2/024001](https://doi.org/10.1088/0264-9381/32/2/024001). arXiv: [1408.3978 \[gr-qc\]](https://arxiv.org/abs/1408.3978).
- [244] Soichiro Morisaki. “Accelerating Parameter Estimation of Gravitational Waves from Compact Binary Coalescence Using Adaptive Frequency Resolutions”. *Phys. Rev. D* 104.4 (2021), p. 044062. DOI: [10.1103/PhysRevD.104.044062](https://doi.org/10.1103/PhysRevD.104.044062). arXiv: [2104.07813 \[gr-qc\]](https://arxiv.org/abs/2104.07813).
- [245] K. Belczynski et al. “The Effect of Pair-Instability Mass Loss on Black Hole Mergers”. *Astron. Astrophys.* 594 (2016), A97. DOI: [10.1051/0004-6361/201628980](https://doi.org/10.1051/0004-6361/201628980). arXiv: [1607.03116 \[astro-ph.HE\]](https://arxiv.org/abs/1607.03116).
- [246] Colm Talbot and Eric Thrane. “Measuring the binary black hole mass spectrum with an astrophysically motivated parameterization”. *Astrophys. J.* 856.2 (2018), p. 173. DOI: [10.3847/1538-4357/aab34c](https://doi.org/10.3847/1538-4357/aab34c). arXiv: [1801.02699 \[astro-ph.HE\]](https://arxiv.org/abs/1801.02699).
- [247] Vassiliki Kalogera et al. “Formation of Double Compact Objects”. *Phys. Rept.* 442 (2007), pp. 75–108. DOI: [10.1016/j.physrep.2007.02.008](https://doi.org/10.1016/j.physrep.2007.02.008). arXiv: [astro-ph/0612144](https://arxiv.org/abs/astro-ph/0612144).

- [248] David W. Hogg. “Distance Measures in Cosmology” (1999). arXiv: [astro-ph/9905116](#).
- [249] P. A. R. Ade et al. “Planck 2015 Results. XIII. Cosmological Parameters”. *Astron. Astrophys.* 594 (2016), A13. DOI: [10.1051/0004-6361/201525830](#). arXiv: [1502.01589 \[astro-ph.CO\]](#).
- [250] Colm Talbot et al. “Parallelized Inference for Gravitational-Wave Astronomy”. *Phys. Rev. D.* 100.4, 043030 (2019), p. 043030. DOI: [10.1103/PhysRevD.100.043030](#). arXiv: [1904.02863 \[astro-ph.IM\]](#).
- [251] R. Abbott et al. “Open Data from the Third Observing Run of LIGO, Virgo, KAGRA and GEO” (2023). arXiv: [2302.03676 \[gr-qc\]](#).
- [252] T. Callister. “A Thesaurus for Common Priors in Gravitational-Wave Astronomy” (2021). arXiv: [2104.09508 \[gr-qc\]](#).
- [253] Tania Regimbau et al. “A Mock Data Challenge for the Einstein Gravitational-Wave Telescope”. *Phys. Rev. D* 86 (2012), p. 122001. DOI: [10.1103/PhysRevD.86.122001](#). arXiv: [1201.3563 \[gr-qc\]](#).
- [254] Benjamin P. Abbott et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. *Phys. Rev. Lett.* 116.6 (2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](#). arXiv: [1602.03837 \[gr-qc\]](#).
- [255] Richard 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 (2021), p. 021053. DOI: [10.1103/PhysRevX.11.021053](#). arXiv: [2010.14527 \[gr-qc\]](#).
- [256] Alexander H. Nitz et al. “4-OGC: Catalog of Gravitational Waves from Compact-Binary Mergers”. 946.2 (2023), p. 59. DOI: [10.3847/1538-4357/aca591](#). arXiv: [2112.06878 \[astro-ph.HE\]](#).

- [257] Tejaswi Venumadhav et al. “A New Search Pipeline for Compact Binary Mergers: Results for Binary Black Holes in the First Observing Run of Advanced LIGO”. *Physical Review D* 100.2 (2019), p. 023011. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.100.023011](https://doi.org/10.1103/PhysRevD.100.023011). arXiv: [1902.10341](https://arxiv.org/abs/1902.10341) [[astro-ph.IM](#)].
- [258] J. Aasi et al. “Advanced LIGO”. *Class. Quant. Grav.* 32 (2015), p. 074001. DOI: [10.1088/0264-9381/32/7/074001](https://doi.org/10.1088/0264-9381/32/7/074001). arXiv: [1411.4547](https://arxiv.org/abs/1411.4547) [[gr-qc](#)].
- [259] Maggie Tse et al. “Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy”. *Phys. Rev. Lett.* 123.23 (2019), p. 231107. DOI: [10.1103/PhysRevLett.123.231107](https://doi.org/10.1103/PhysRevLett.123.231107).
- [260] Fausto Acernese et al. “Increasing the Astrophysical Reach of the Advanced Virgo Detector via the Application of Squeezed Vacuum States of Light”. *Phys. Rev. Lett.* 123.23 (2019), p. 231108. DOI: [10.1103/PhysRevLett.123.231108](https://doi.org/10.1103/PhysRevLett.123.231108).
- [261] Benjamin P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. *Phys. Rev. Lett.* 119.16 (2017), p. 161101. DOI: [10.1103/PhysRevLett.119.161101](https://doi.org/10.1103/PhysRevLett.119.161101). arXiv: [1710.05832](https://arxiv.org/abs/1710.05832) [[gr-qc](#)].
- [262] Benjamin P. Abbott et al. “Multi-Messenger Observations of a Binary Neutron Star Merger”. *Astrophys. J. Lett.* 848.2 (2017), p. L12. DOI: [10.3847/2041-8213/aa91c9](https://doi.org/10.3847/2041-8213/aa91c9). arXiv: [1710.05833](https://arxiv.org/abs/1710.05833) [[astro-ph.HE](#)].
- [263] M. Campanelli, C. O. Lousto, and Y. Zlochower. “Spinning-Black-Hole Binaries: The Orbital Hang-Up”. *Physical Review D* 74.4 (2006), p. 041501. DOI: [10.1103/PhysRevD.74.041501](https://doi.org/10.1103/PhysRevD.74.041501). arXiv: [gr-qc/0604012](https://arxiv.org/abs/gr-qc/0604012).
- [264] Emily Baird et al. “Degeneracy between Mass and Spin in Black-Hole-Binary Waveforms”. *Physical Review D* 87.2 (2013), p. 024035. DOI: [10.1103/PhysRevD.87.024035](https://doi.org/10.1103/PhysRevD.87.024035). arXiv: [1211.0546](https://arxiv.org/abs/1211.0546) [[gr-qc](#)].
- [265] Sascha Husa et al. “Frequency-Domain Gravitational Waves from Nonprecessing Black-Hole Binaries. I. New Numerical Waveforms and Anatomy of the Signal”.

- Physical Review D* 93.4 (2016), p. 044006. DOI: [10.1103/PhysRevD.93.044006](https://doi.org/10.1103/PhysRevD.93.044006). arXiv: [1508.07250](https://arxiv.org/abs/1508.07250) [gr-qc].
- [266] Sebastian Khan et al. “Frequency-Domain Gravitational Waves from Nonprecessing Black-Hole Binaries. II. A Phenomenological Model for the Advanced Detector Era”. *Phys. Rev. D* 93.4 (2016), p. 044007. DOI: [10.1103/PhysRevD.93.044007](https://doi.org/10.1103/PhysRevD.93.044007). arXiv: [1508.07253](https://arxiv.org/abs/1508.07253) [gr-qc].
- [267] Andrea Taracchini et al. “Prototype Effective-One-Body Model for Nonprecessing Spinning Inspirational-Merger-Ringdown Waveforms”. *Physical Review D* 86.2 (2012), p. 024011. DOI: [10.1103/PhysRevD.86.024011](https://doi.org/10.1103/PhysRevD.86.024011). arXiv: [1202.0790](https://arxiv.org/abs/1202.0790) [gr-qc].
- [268] Andrea Taracchini et al. “Effective-One-Body Model for Black-Hole Binaries with Generic Mass Ratios and Spins”. *Physical Review D* 89.6 (2014), p. 061502. DOI: [10.1103/PhysRevD.89.061502](https://doi.org/10.1103/PhysRevD.89.061502). arXiv: [1311.2544](https://arxiv.org/abs/1311.2544) [gr-qc].
- [269] Sebastian Khan et al. “Phenomenological Model for the Gravitational-Wave Signal from Precessing Binary Black Holes with Two-Spin Effects”. *Physical Review D* 100.2 (2019), p. 024059. DOI: [10.1103/PhysRevD.100.024059](https://doi.org/10.1103/PhysRevD.100.024059). arXiv: [1809.10113](https://arxiv.org/abs/1809.10113) [gr-qc].
- [270] Héctor Estellés et al. “Phenomenological Time Domain Model for Dominant Quadrupole Gravitational Wave Signal of Coalescing Binary Black Holes”. *Phys. Rev. D* 103.12 (2021), p. 124060. DOI: [10.1103/PhysRevD.103.124060](https://doi.org/10.1103/PhysRevD.103.124060). arXiv: [2004.08302](https://arxiv.org/abs/2004.08302) [gr-qc].
- [271] Antoni Ramos-Buades et al. *SEOBNRv5PHM: Next Generation of Accurate and Efficient Multipolar Precessing-Spin Effective-One-Body Waveforms for Binary Black Holes*. 2023. URL: <https://dcc.ligo.org/P2300067-v4>.
- [272] Eleanor Hamilton et al. “Model of Gravitational Waves from Precessing Black-Hole Binaries through Merger and Ringdown”. *Physical Review D* 104.12 (2021), p. 124027. DOI: [10.1103/PhysRevD.104.124027](https://doi.org/10.1103/PhysRevD.104.124027). arXiv: [2107.08876](https://arxiv.org/abs/2107.08876) [gr-qc].

- [273] Alessandra Buonanno and Thibault Damour. “Transition from Inspiral to Plunge in Binary Black Hole Coalescences”. *Physical Review D* 62.6 (2000), p. 064015. DOI: [10.1103/PhysRevD.62.064015](https://doi.org/10.1103/PhysRevD.62.064015). arXiv: [gr-qc/0001013](https://arxiv.org/abs/gr-qc/0001013).
- [274] Thibault Damour, Piotr Jaranowski, and Gerhard Schäfer. “Determination of the Last Stable Orbit for Circular General Relativistic Binaries at the Third Post-Newtonian Approximation”. *Physical Review D* 62.8 (2000), p. 084011. DOI: [10.1103/PhysRevD.62.084011](https://doi.org/10.1103/PhysRevD.62.084011). arXiv: [gr-qc/0005034](https://arxiv.org/abs/gr-qc/0005034).
- [275] Thibault Damour. “Coalescence of Two Spinning Black Holes: An Effective One-Body Approach”. *Physical Review D* 64.12 (2001), p. 124013. DOI: [10.1103/PhysRevD.64.124013](https://doi.org/10.1103/PhysRevD.64.124013). arXiv: [gr-qc/0103018](https://arxiv.org/abs/gr-qc/0103018).
- [276] Alessandra Buonanno, Yanbei Chen, and Thibault Damour. “Transition from Inspiral to Plunge in Precessing Binaries of Spinning Black Holes”. *Physical Review D* 74.10 (2006), p. 104005. DOI: [10.1103/PhysRevD.74.104005](https://doi.org/10.1103/PhysRevD.74.104005). arXiv: [gr-qc/0508067](https://arxiv.org/abs/gr-qc/0508067).
- [277] Lorenzo Pompili et al. “Laying the Foundation of the Effective-One-Body Waveform Models SEOBNRv5: Improved Accuracy and Efficiency for Spinning Non-Precessing Binary Black Holes” (2023). DOI: [10.48550/arXiv.2303.18039](https://doi.org/10.48550/arXiv.2303.18039). arXiv: [2303.18039](https://arxiv.org/abs/2303.18039) [[gr-qc](https://arxiv.org/abs/gr-qc)].
- [278] Alessandro Nagar et al. “Multipolar Effective One Body Waveform Model for Spin-Aligned Black Hole Binaries”. *Physical Review D* 102.2 (2020), p. 024077. DOI: [10.1103/PhysRevD.102.024077](https://doi.org/10.1103/PhysRevD.102.024077). arXiv: [2001.09082](https://arxiv.org/abs/2001.09082) [[gr-qc](https://arxiv.org/abs/gr-qc)].
- [279] Thibault Damour and Alessandro Nagar. “Improved Analytical Description of Inspiralling and Coalescing Black-Hole Binaries”. *Physical Review D* 79.8 (2009), p. 081503. DOI: [10.1103/PhysRevD.79.081503](https://doi.org/10.1103/PhysRevD.79.081503). arXiv: [0902.0136](https://arxiv.org/abs/0902.0136) [[gr-qc](https://arxiv.org/abs/gr-qc)].
- [280] Thibault Damour and Alessandro Nagar. “New Effective-One-Body Description of Coalescing Nonprecessing Spinning Black-Hole Binaries”. *Physical Review D* 90.4 (2014), p. 044018. DOI: [10.1103/PhysRevD.90.044018](https://doi.org/10.1103/PhysRevD.90.044018). arXiv: [1406.6913](https://arxiv.org/abs/1406.6913) [[gr-qc](https://arxiv.org/abs/gr-qc)].

- [281] Alessandro Nagar et al. “Energetics and Phasing of Nonprecessing Spinning Coalescing Black Hole Binaries”. *Physical Review D* 93.4 (2016), p. 044046. DOI: [10.1103/PhysRevD.93.044046](https://doi.org/10.1103/PhysRevD.93.044046). arXiv: [1506.08457](https://arxiv.org/abs/1506.08457) [gr-qc].
- [282] Geraint Pratten et al. “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 (2020), p. 064001. DOI: [10.1103/PhysRevD.102.064001](https://doi.org/10.1103/PhysRevD.102.064001). arXiv: [2001.11412](https://arxiv.org/abs/2001.11412) [gr-qc].
- [283] Cecilio García-Quirós et al. “Multimode Frequency-Domain Model for the Gravitational Wave Signal from Nonprecessing Black-Hole Binaries”. *Phys. Rev. D* 102.6 (2020), p. 064002. DOI: [10.1103/PhysRevD.102.064002](https://doi.org/10.1103/PhysRevD.102.064002). arXiv: [2001.10914](https://arxiv.org/abs/2001.10914) [gr-qc].
- [284] Héctor Estellés et al. “Time-Domain Phenomenological Model of Gravitational-Wave Subdominant Harmonics for Quasicircular Nonprecessing Binary Black Hole Coalescences”. *Phys. Rev. D* 105.8 (2022), p. 084039. DOI: [10.1103/PhysRevD.105.084039](https://doi.org/10.1103/PhysRevD.105.084039). arXiv: [2012.11923](https://arxiv.org/abs/2012.11923) [gr-qc].
- [285] Héctor Estellés et al. “New Twists in Compact Binary Waveform Modeling: A Fast Time-Domain Model for Precession”. *Phys. Rev. D* 105.8 (2022), p. 084040. DOI: [10.1103/PhysRevD.105.084040](https://doi.org/10.1103/PhysRevD.105.084040). arXiv: [2105.05872](https://arxiv.org/abs/2105.05872) [gr-qc].
- [286] Tousif Islam et al. “Surrogate Model for Gravitational Wave Signals from Non-Spinning, Comparable- to Large-Mass-Ratio Black Hole Binaries Built on Black Hole Perturbation Theory Waveforms Calibrated to Numerical Relativity”. 106.10 (2022), p. 104025. DOI: [10.1103/PhysRevD.106.104025](https://doi.org/10.1103/PhysRevD.106.104025). arXiv: [2204.01972](https://arxiv.org/abs/2204.01972) [gr-qc].
- [287] M. Punturo et al. “The Third Generation of Gravitational Wave Observatories and Their Science Reach”. *Classical and Quantum Gravity* 27.8 (2010), p. 084007. ISSN: 0264-9381. DOI: [10.1088/0264-9381/27/8/084007](https://doi.org/10.1088/0264-9381/27/8/084007).

BIBLIOGRAPHY

- [288] Matthew Evans et al. “A Horizon Study for Cosmic Explorer: Science, Observatories, and Community” (2021). DOI: [10.48550/arXiv.2109.09882](https://doi.org/10.48550/arXiv.2109.09882). arXiv: [2109.09882](https://arxiv.org/abs/2109.09882) [[astro-ph.IM](https://arxiv.org/abs/2109.09882)].