
**SEARCHING FOR PRECESSING BLACK HOLE BINARIES IN
GRAVITATIONAL-WAVE DATA**

Stefano Schmidt

ISBN: 978-94-6496-395-3

DOI: <https://doi.org/10.33540/2857>

Printed by: Gildeprint – www.gildeprint.nl

Copyright: © 2025 Stefano Schmidt



This work originates as part of the research program of the Foundation for Fundamental Research on Matter (FOM), and falls as of April 1, 2017 under the responsibility of the Foundation for Nederlandse Wetenschappelijk Onderzoek Instituten (NWO-I), which is part of the Dutch Research Council (NWO).

Searching for Precessing Black Hole Binaries in Gravitational-wave Data

**Zoeken naar Dubbele Zwarte Gaten met Precessie in
Zwaartekrachtsgolven-data**

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de
Universiteit Utrecht
op gezag van de
rector magnificus, prof. dr. ir. W. Hazleger,
ingevolge het besluit van het College voor Promoties
in het openbaar te verdedigen op
dinsdag 20 mei 2025 des middags te 12.15 uur

door

Stefano Schmidt
geboren op 21 juli 1995
te MILAAN, Italië

Promotor:

Prof. dr. C.F.F. van den Broeck

Copromotor:

Dr. S.E. Caudill

Beoordelingscommissie:

Prof. dr. J.F.J. van der Brand

Prof. dr. S. Fairhurst

Dr. A. Grelli

Prof. dr. T. Peitzmann

Dr. A. Samajdar

Vanitas vanitatum, dixit Ecclesiastes; vanitas vanitatum, et omnia vanitas.

Quid habet amplius homo de universo labore suo quo laborat sub sole?

—Qohelet, 1:2-3

Contents

Index	viii
Preface	ix
1 Basic Concepts about Gravitational Waves	1
1.1 Gravitational Waves in General Relativity	1
1.1.1 Essentials of General Relativity	1
1.1.2 Linearized Gravity	6
1.1.3 Physical Degrees of Freedom of a Gravitational Wave	7
1.1.4 Interaction with Test Masses	9
1.1.5 Energy of a Gravitational Wave	11
1.2 Sources of Gravitational Waves	12
1.2.1 The Quadrupole Formula	13
1.2.2 Multipole Expansion of the Source and of $h(t)$	14
1.3 Black Hole Binaries	17
1.3.1 Black Hole and Black Hole Binaries in Astrophysics	17
1.3.2 Orbital Evolution of a Black Hole Binary	18
1.3.3 Precession in Black Hole Binaries	24
2 Detecting Gravitational Waves with Data Analysis	31
2.1 Interferometers for Gravitational-wave Detection	31
2.1.1 Features of a Real Interferometer	35
2.1.2 Sources of Noise	36
2.2 Noise Characterization	39
2.2.1 Essentials of Timeseries Analysis	39
2.2.2 A statistical Model for the Noise	40
2.2.3 Spectral Estimation	44
2.2.4 Whitening	45

2.3	Overview of the Data Analysis for Compact Binary Coalescences	46
3	Searching for Signals with Matched Filtering	51
3.1	Basics of Matched Filtering	51
3.1.1	Matched Filtering as Optimal Filter	52
3.1.2	Matched Filtering as Hypothesis Test	53
3.2	Search Statistic	54
3.2.1	Statistic for Circular Aligned-spin Binaries	56
3.2.2	Statistic for General Binaries	58
3.2.3	Maximization over Time	60
3.3	Template Banks	62
3.4	Features of a Realistic Pipeline	66
3.4.1	Generating Triggers	67
3.4.2	Signal Consistency Test	69
3.4.3	Ranking	70
3.4.4	False Alarm Rate	71
3.4.5	<i>p</i> astro	72
4	Generating Template Banks in High Dimensional Spaces	75
4.1	Introduction	75
4.2	Methods	77
4.2.1	The Metric	78
4.2.2	Sampling from the Manifold	80
4.2.3	Random Template Placing	83
4.3	Validation	85
4.3.1	Normalizing Flow Validation	87
4.3.2	Template Placement Performance	88
4.4	Comparison with Other Bank Generation Methods	89
4.4.1	A Non-spinning HM Template Bank	90
4.4.2	An “All-sky” Template Bank	91
4.5	Novel Applications of the Method	93
4.5.1	A Precessing Bank	93
4.5.2	An Aligned-spin HM Bank	100
4.5.3	Other Possible Applications	103
4.6	Future Prospects	106
4.7	Final Remarks	109
4.A	Details of the Metric Computation	110
4.B	Alternative Definitions for the Metric	112
4.C	Computing the Volume of the Parameter Space	114

5 Generalized Signal Consistency Test	115
5.1 Introduction	115
5.2 A New Generalized ξ^2 Signal consistency Test	116
5.2.1 The “Standard” Signal Consistency Test	117
5.2.2 The Novel “Symphony” Signal Consistency Test	118
5.2.3 Approximating the New ξ^2 Test	120
5.3 Validity and Limitations of Different Signal Consistency Tests	123
5.3.1 How Does the Different Tests Compare to Each Other?	124
5.3.2 When Does the “Standard” Test Fail?	127
5.3.3 How Does the “Mixed” Test Perform?	127
5.3.4 Which “Mixed” Test Should We Use?	130
5.3.5 How to Choose the Autocorrelation Length for the Test?	131
5.3.6 How Does Real Noise Affect the Test Performance?	132
5.4 Final Remarks	133
5.A Expected Value of the “Symphony” Signal Consistency Test in Gaussian Noise	134
6 Searching for Gravitational-wave Signals from Precessing Binaries	137
6.1 Introduction	137
6.2 Search Method	139
6.3 Template Banks	141
6.3.1 SVD Compression	148
6.4 Precessing Searches Results	150
6.4.1 Recovered SNR	150
6.4.2 ξ^2 Test Results	152
6.4.3 Sensitivity Improvement	154
6.5 Outline of Future Improvements	157
6.6 Final Remarks	159
7 Searching for Precessing Binaries in LIGO Data from the Third Observing Run	161
7.1 Introduction	161
7.2 Set-up of the Search	162
7.3 Search Results	163
7.3.1 Rate Upper Limits	167
7.4 Conclusion	169
8 Closing Remarks	171
Public Summary	175

Openbare Samenvatting	181
Selected Publications	187
List of Figures	198
List of Tables	200
Acknowledgments	201
Bibliography	241

Preface

The theory of General Relativity, proposed by Albert Einstein in 1915, provides today the most accurate description of the gravitational interaction. The theory models the spacetime as a geometric manifold, whose structure is influenced by the distribution of masses across the spacetime. In turn, the geometry of the spacetime influences the movement of free falling masses, hence affecting the mass distribution itself. The theory has been successfully applied to the a wide variety of physical phenomena on different scales, ranging from the motion of the planets to the cosmological evolution of the universe. Among many others phenomena, general relativity predicts the existence of very dense star-like objects, called *black holes*, and of tiny propagating high-frequency perturbations to the “slowly moving” spacetime structure, called *gravitational waves*. They will be both of primary interest for our work.

A black hole is a region of spacetime where the effect of gravity is so strong that nothing can posses enough speed to escape. Clearly, such strong gravity must be generated by some high density mass distribution and for this reason, black holes are thought to be the results of the gravitational collapse of ordinary matter, triggered whenever the outward pressure generated by matter is not able to support the inward pressure of gravity. By the aptly named *no-hair* theorem, the effect of a black hole on the spacetime is uniquely defined by its mass, spin and electric charge [10].

Gravitational waves are the propagating effect on spacetime of a time-varying mass distribution. They arise naturally from causality, which requires the gravitational interaction to propagate at finite speed. Gravitational waves have the effect of changing the distance between two free-falling observers and such variation can be measured by sophisticated interferometric-based detectors, such as LIGO [11], Virgo [12] or KAGRA [13], giving birth to the field of gravitational-wave astronomy [14].

Due to their time-varying nature, black hole binaries are copious emitters of gravitational waves and indeed, more than 90 black hole binaries were detected so far by the LIGO-Virgo-KAGRA collaboration [15–18], thanks to their gravitational wave

emission. The significance of such observations cannot be emphasized enough. The direct detection of a gravitational wave from black hole binaries allowed for advancements in astrophysics [19–21], fundamental physics [22–25], cosmology [26–29] and general relativity [30–34], with promises of even more groundbreaking impacts with the next generation observatories, such as the Einstein Telescope [35,36], Cosmic Explorer [37] and LISA [38].

According to the general relativistic description of a binary black hole, the two black hole spins play an important role in determining the dynamics of the binary. Indeed, their magnitude and mutual interaction can speed up or slow down the gravitational wave emission and hence the rate of orbital shrinking. Moreover, if at least one of the two spins is mis-aligned with the orbital angular momentum [39], the interaction between spin and orbital angular momentum \mathbf{L} causes the orbital plane to rotate around an approximately constant axis. This phenomenon is called *precession* and the gravitational wave signal emitted in this case acquires a more complicated structure [40–43].

Despite the intricacies inherent in the modeling [44–52], the underlying physical concept is clear: as gravitational wave emission attains its maximum along the direction $\hat{\mathbf{L}}$ of the orbital angular momentum, and given that $\hat{\mathbf{L}}$ evolves over time, an inertial observer will measure a time modulation of the wave amplitude and phase, as the orbital plane points towards and away from the observer. This effect can be observed by current gravitational wave detectors and indeed, precession has been observed as a statistical property of the BBH population [20,21]: this marks yet another remarkable achievement of the LIGO, Virgo and KAGRA collaboration. However, very few individual signals show conclusive evidence for largely misaligned spins [53], interesting exceptions including GW190521 [54], GW191109 [55] and GW200129 [56].

The primary scientific interest of precession arises from the fact that through a single observation of a heavily precessing binary we can break the degeneracies in the measurement of several parameters that characterize the binary, thus achieving higher accuracy in the inferred physical quantities. For example, breaking the degeneration between distance and inclination has a direct impact on the accuracy of cosmological parameters inferred using gravitational wave observations [27,57]. As another example, the measurement of individual spins enables a better understanding of the spin distribution in the population of black hole binaries, which directly impacts our understanding of the binaries' formation mechanism [58–62]¹.

The binary black hole detections made so far have been achieved through mod-

¹Indeed, population studies indicate that heavily precessing systems are more likely to form through dynamic assembly rather than through common evolution [61]. However, currently the fraction of binaries formed from each channel and a detailed understanding of these channels are still under debate and could possibly be illuminated by further detections of precessing signals.

eled searches conducted by advanced matched-filtering pipelines. These pipelines use matched filtering [63] to correlate the detector’s output with a large number of waveform templates, gathered in so-called *template banks*. While modeled searches are optimal for known signals in purely Gaussian noise, they rapidly loose sensitivity for signals not included among the templates [64–68]. Unfortunately, despite having possibly detected a few precessing signals, the modeled searches deployed to build the modern catalogs [15–18, 69–72] search only for “aligned-spin” systems (where the black holes spins are aligned or anti-aligned with the orbital angular momentum) and do not include the effects of precession in the models used for the template waveforms. Indeed, as aligned-spin systems have only two spin-related degrees of freedom² and result in a simpler waveform structure, building a search for gravitational waves from merging aligned-spin binary black holes drastically simplifies the search problem and reduces the computational cost, hence making a systematic search of systems over a broad mass range feasible. On the other hand, only deploying aligned-spin searches limit the sensitivity of current pipelines towards precessing signals in certain regions of the parameter space, most notably for asymmetric systems and maximally mis-aligned spins [3, 73, 74]. For instance, as pointed out in [75], an aligned-spin search targeting neutron star-black hole systems can miss up to $\sim 60\%$ of highly precessing sources with mass ratio $q \gtrsim 6$ and largely mis-aligned spins.

The limited sensitivity of traditional matched-filter pipelines towards precessing systems poses an important question, which we address in this dissertation. Is the observed lack of strongly precessing signals due to their rarity [76]? Or, rather, is it caused by the limited sensitivity of current searches to such extreme signals? To discern between the two alternatives, it is important to develop a search tuned for heavily precessing systems and able to achieve optimal sensitivity, thus unlocking the scientific potential of studying precession in black hole binaries: this is the goal of the research described here. But the importance of developing a search for precessing signals extends beyond the question posed above. Indeed, the machinery we develop here can be straightforwardly applied to matched-filtering searches of other types of signals (e.g. from eccentric binaries), hence paving the way to a large variety of studies in gravitational-wave astronomy: these future prospects will be further discussed in the conclusions of the dissertation.

Moving from an aligned-spin search to a precessing search is a highly non trivial task, which requires us to address a number of challenges:

- *Increased number of degrees of freedom for the search templates.* The large variability of precessing signals makes the current template placement techniques un-

²As opposed to six spin-related degrees of freedom of a precessing binary black hole.

feasibly computationally costly and urges for novel techniques to be employed. In [1], we addressed this by developing a novel machine learning based method for template bank generation, particularly suited for this scenario.

- *Generalized signal consistency test.* In order to distinguish between candidates from astrophysical and terrestrial origin, each search performs a consistency test between the observed data and the template signals. In the precessing case, the signal consistency test needs to be heavily modified and its validity assessed: this is the main contribution we gave in [2].
- *Choice of a suitable region of the parameter space to search.* Due to the large number of templates required to carry on a precessing search, it is crucial to select a range in the binary masses and spin where to focus the search on. Ideally, a precessing search on the target binaries should maximise the sensitivity improvement over a traditional search. In [3], we identified a target region of the parameter space where such improvement is at hand and we performed a proof-of-principle search to demonstrate the sensitivity increase.

In this dissertation, we describe the long journey toward searching for precessing signals using the GstLAL pipeline [77–82]. In Chapter 1 we begin with an introduction to gravitational waves and their detection principles and to black hole binaries (with a particular focus on precession). Chapter 2 is devoted to the discussion of basic data analysis techniques for gravitational wave astronomy, while in Chapter 3 we delve into the details of the matched-filtering searches for gravitational waves from binary black hole signals. The remaining chapters address the challenges outlined above for searches for precessing signals. In Chapter 4 we discuss a novel template placement method, especially designed for complex signals such as those associated to precession. Chapter 5 introduces a generalized signal consistency test and discusses its limitations and applicability to precessing searches. In Chapter 6 we present a study on the performance of two precessing searches, covering two different regions of the parameter space, and we assess the sensitivity improvement enabled by our search method. Finally, in Chapter 7 we look for precessing signals in the publicly available data from the third observing run [83, 84] from the LIGO-Virgo-KAGRA collaboration. While we do not report new detections, our results allow us to place an upper limit on the rate of a potential *undetected* population of precessing signals. Chapter 8 summarizes our work and gathers some closing remarks, highlighting possible future developments of this work.

CHAPTER 1

Basic Concepts about Gravitational Waves

This chapter provides a concise overview on the field of gravitational waves (GWs). It is by no means meant to be exhaustive but it acts as a summary of the physical concepts needed to understand our work. After providing an introduction to gravitational waves in the context of the theory of general relativity and to its possible sources, we will discuss their effect on matter and their detection principles. Finally, we will discuss with some details the importance of black holes binaries (BBH) and their dynamics, with a special attention to the phenomenon of precession.

Throughout the chapter, we will closely follow the discussion in [85], and unless explicitly noted, we employ geometrical units, where $G = c = 1$.

1.1 Gravitational Waves in General Relativity

1.1.1 Essentials of General Relativity

General relativity (GR) is the theory that represents our most advanced knowledge of gravity and at its core, the theory describes the way that masses (or equivalently an energy distribution) affect each others' motion [86]. Despite the intricacies of the mathematics, the general idea is simple: the spacetime has a geometric structure that defines a set of curves, called *geodesics*, which correspond to the motion of free particles, i.e. particles not subject to any force. The geometric structure of the spacetime, which influences the motion of the masses, is affected by the masses themselves. In this way, the masses filling the spacetime interact with each other in a non-linear and

non-trivial way, giving rise to the complex phenomenology of GR.

The mathematical formulation of the theory heavily relies on differential geometry to describe the spacetime. According to this mathematical model, a generic point of the space time, also called *event*, is represented as a four-vector x^μ on a 4 dimensional semi-Riemannian manifold. For example, a straightforward choice of coordinates can label an event using a time coordinate t and a Cartesian space coordinate \mathbf{x} , so that:

$$x^\mu = (t, \mathbf{x}). \quad (1.1)$$

Clearly, the choice of coordinate system is arbitrary and it can be changed through a *coordinate transformation* mapping the old coordinates x^μ to the new ones x'^μ :

$$x^\mu \longrightarrow x'^\mu \quad (1.2)$$

where the ' symbol will be used to denote the new coordinate system. As any physical law or prediction cannot depend on the chosen coordinate frame (*general covariance*), a global transformation of the spacetime coordinate is the *gauge symmetry* of the theory of general relativity.

According to GR, all the physical quantities are represented by tensors. Mathematically, a tensor of type $(k - l)$ is an element of a suitable vector space of dimension $k + l^1$. For the purpose of our discussion, it is sufficient to state that, given a *canonical* choice of basis, each tensor can be represented by its components, which are multi-index quantity $T_{\nu_1 \dots \nu_l}^{\mu_1 \dots \mu_k}$. By defining a tensor that varies smoothly at each point of the spacetime manifold, we create a tensor field $T_{\nu_1 \dots \nu_l}^{\mu_1 \dots \mu_k}(x^\mu)$, which we will still call tensor, with a slight abuse of notation. A defining property of tensors is their behaviour under general coordinate transformation. More precisely, under the change of coordinates Eq. (1.2) a $(k - l)$ tensor $T_{\nu_1 \dots \nu_l}^{\mu_1 \dots \mu_k}$ transforms as

$$T_{\nu_1 \dots \nu_l}^{\mu_1 \dots \mu_k}(x^\mu) \longrightarrow T'_{\nu_1 \dots \nu_l}^{\mu_1 \dots \mu_k}(x'^\mu) = T_{\sigma_1 \dots \sigma_l}^{\rho_1 \dots \rho_k} \frac{\partial x'^{\mu_1}}{\partial x^{\rho_1}} \dots \frac{\partial x'^{\mu_k}}{\partial x^{\rho_k}} \frac{\partial x^{\sigma_1}}{\partial x'^{\nu_1}} \dots \frac{\partial x^{\sigma_l}}{\partial x'^{\nu_l}} \quad (1.3)$$

where we employed the Einstein summation convention for repeated indices:

$$a^\mu b_\mu = \sum_\mu a^\mu b_\mu. \quad (1.4)$$

According to an equivalent characterization, a tensor of type $(k - l)$ can also be seen as a linear map (a functional) between tensor of type $(l - k)$ and a scalar, i.e. a real quantity invariant under a coordinate transformation. For example, a $(1 - 2)$ tensor

¹More formally, a $(k - l)$ tensor is defined as a linear functional acting on k -vectors and l -forms. Thanks to the linearity of its action, it can be shown that all the set of all the $(k - l)$ tensors satisfies the properties of a vector space, hence the definition given above.

$U_{\mu\nu}^\rho$ can be applied to a tensor $V_\rho^{\mu\nu}$ of type $(2 - 1)$ to yield the scalar $s(x^\mu)$, by an operation called *tensor contraction*:

$$s(x^\mu) = U_{\mu\nu}^\rho V_\rho^{\mu\nu}. \quad (1.5)$$

By contracting together two tensors, we can define other tensors. For instance by combining a $(3 - 2)$ tensor $U_{\mu\nu}^{\rho\sigma\tau}$ with a $(0 - 2)$ tensor $V_{\alpha\beta}$, we can obtain a $(1 - 2)$ tensor Z with components $Z_{\mu\nu}^\rho = U_{\mu\nu}^{\rho\alpha\beta} V_{\alpha\beta}$. We may also generate new tensors by applying the tensor product, denoted by \otimes , between two tensors A, B :

$$(A \otimes B)_{\nu_1 \dots \nu_l \sigma_1 \dots \sigma_l}^{\mu_1 \dots \mu_k \rho_1 \dots \rho_k} = A_{\nu_1 \dots \nu_l}^{\mu_1 \dots \mu_k} B_{\sigma_1 \dots \sigma_l}^{\rho_1 \dots \rho_k}. \quad (1.6)$$

The advantage of using tensors in GR relies on the fact that the tensor transformation laws Eq. (1.3) under general coordinate transformations make straightforward to construct scalar quantities satisfying the gauge symmetry of general relativity by contraction, as shown in Eq. (1.5). Indeed by using $\frac{\partial x'^\mu}{\partial x^\alpha} \frac{\partial x^\alpha}{\partial x'^\nu} = \delta_\nu^\mu$, it is a simple exercise, to show that any scalar produced by tensor contraction is invariant under change of coordinates:

$$s(x^\mu) = s'(x'^\mu). \quad (1.7)$$

This is the reason why we require that any physical quantity in GR is represented by a tensor and we will only formulate equations by relating the components of tensor fields defined on the spacetime. This ensures that the general covariance of the GR is respected.

A special tensor is the metric tensor $g_{\mu\nu}(x^\alpha)$, which can be used to compute the distance ds^2 between two close-by events x_1^μ, x_2^μ , separated by a small displacement dx^μ :

$$ds^2 = g_{\mu\nu}(x^\alpha) dx^\mu dx^\nu. \quad (1.8)$$

The functional ds^2 is also called line element and can be used to compute the length L of a given path $x^\mu(t)$ on the spacetime manifold:

$$L = \int dt \sqrt{g_{\mu\nu}(x^\alpha(t)) \frac{dx^\mu}{dt} \frac{dx^\nu}{dt}}. \quad (1.9)$$

We also introduce the inverse metric tensor $g^{\mu\nu}$ such that $g^{\mu\rho} g_{\nu\rho}$ is equal to the identity δ_ν^μ .

A spacetime is called *flat* if the metric tensor is $g^{\mu\nu} = \eta^{\mu\nu} = \text{diag}(-1, 1, 1, 1)$. Due to non flat metric tensor, the partial derivative operator $\partial_\alpha = \frac{\partial}{\partial x^\alpha}$ is not a tensor, since it does not satisfy the tensor transformation law Eq. (1.3). To overcome this shortcoming, we introduce the covariant derivative operator ∇_μ , acting on a vector v^μ as:

$$\nabla_\mu v^\nu = \partial_\mu v^\nu + \Gamma_{\mu\sigma}^\nu v^\sigma \quad (1.10)$$

where the quantities $\Gamma_{\rho\sigma}^\mu$ are called Christoffel symbols and they are functions of the metric and its derivatives

$$\Gamma_{\rho\sigma}^\mu = \frac{1}{2} g^{\mu\alpha} (\partial_\rho g_{\sigma\alpha} + \partial_\sigma g_{\rho\alpha} - \partial_\alpha g_{\rho\sigma}). \quad (1.11)$$

Note that if the metric is flat, i.e. $\partial_\rho g_{\sigma\alpha} = 0$, the Christoffel symbols vanish and $\nabla_\mu = \partial_\mu$.

The metric tensor is important not only for computing distances and defining derivatives but also for determining the trajectory of a test particle—a particle with negligible mass and no external forces acting upon it—in spacetime. In this context, the metric tensor $g_{\mu\nu}$ plays the role of the gravitational potential in Newtonian gravity. More precisely, general relativity² postulates that a free particle follows a *geodesic*, which is a “straight” line minimizing the distance between its two endpoints, and that the shape of a geodesic is fully specified by the metric tensor. To provide a characterization for a geodesics, we need to consider that, like in the Newtonian case, the velocity $v^\mu = \frac{dx^\mu}{dt}$ of a geodesics path $x^\mu(t)$ is constant along the path itself. Mathematically, this amounts to the requirement that the directional covariant derivative of v^μ along the path (i.e. along the direction of v^μ) is zero:

$$v^\nu \nabla_\nu v^\mu = 0. \quad (1.12)$$

Note that the equation depends on the metric tensor through the Christoffel symbols appearing in the ∇^μ operator. The equation is called *geodesic equation* and leads to the following equation for the trajectory of a test particle $x^\mu(t)$:

$$\frac{d^2 x^\mu(t)}{dt^2} = \Gamma_{\rho\sigma}^\mu \frac{dx^\rho}{dt} \frac{dx^\sigma}{dt}. \quad (1.13)$$

This is the equation of motion of a particle within the GR framework and it is the general relativistic generalization of the second Newton equation. Indeed, it can be shown that, in the low density, low velocity regime, Eq. (1.13) reduces to the Newtonian equation of motion.

As the motion of particles is fully specified by the metric tensor $g_{\mu\nu}$ and by the energy content $T^{\mu\nu}$ of the spacetime, the theory must provide an equation connecting those two quantities. Following the spirit of the theory, the latter must be represented by a tensor, called *energy momentum tensor* $T^{\mu\nu}$, whose components are defined in terms of the energy and momentum density measured at each point of the spacetime. More precisely, the T^{00} component amounts to the energy density in the comoving frame, while the T^{0i} components amount to the flux of energy across the surface i and the T^{ij} is the stress matrix.

²Indeed the same postulate is also in the Newton's first law of motion.

The relation between the energy momentum tensor and the metric tensor is elegantly established in the famous Einstein equations:

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi T_{\mu\nu} \quad (1.14)$$

where $R_{\mu\nu}$ is the Ricci tensor and $R = R_{\mu}^{\mu} = g^{\mu\nu}R_{\mu\nu}$ is the curvature scalar. The tensor $R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R$ is sometime also referred to as *Einstein tensor* $G_{\mu\nu}$. The Ricci tensor is obtained by contracting two indices of the Riemann curvature tensor $R_{\mu\sigma\nu}^{\rho}$:

$$R_{\mu\nu} = R_{\mu\sigma\nu}^{\sigma} = g^{\sigma\gamma}g_{\rho\gamma}R_{\mu\sigma\nu}^{\rho}. \quad (1.15)$$

The Riemann tensor is a primary measure of the spacetime curvature and, heuristically, it quantifies the failure of two infinitesimally close geodesics, initially parallel to each other, to follow parallel paths along their trajectory. Let t^{μ} be the tangent vector and x^{μ} their separation vector, two observers along the two geodesics will then experience a relative acceleration a^{μ} proportional to the Riemann tensor³:

$$a^{\rho} = R_{\mu\sigma\nu}^{\rho}x^{\sigma}t^{\mu}t^{\nu}. \quad (1.16)$$

This expression is also known as *geodesics deviation equation*. The Riemann tensor can be expressed as a function of the metric tensor, through the derivatives of the Christoffel symbols:

$$R_{\mu\sigma\nu}^{\rho} = \partial_{\sigma}\Gamma_{\mu\nu}^{\rho} - \partial_{\nu}\Gamma_{\sigma\mu}^{\rho} + \Gamma_{\sigma\lambda}^{\rho}\Gamma_{\mu\nu}^{\lambda} - \Gamma_{\nu\lambda}^{\rho}\Gamma_{\mu\sigma}^{\lambda}. \quad (1.17)$$

This expression, along with the definition of the Christoffel symbols Eq. (1.11) and the Einstein equations Eq. (1.14), constitutes a set of coupled partial differential equations for the metric tensor $g^{\mu\nu}$. These equations fully describe the dynamics of spacetime once the energy content $T^{\mu\nu}$ is specified.

We note that the Einstein equations implies the conservation law for the energy momentum tensor. Indeed, as the following identity is true for any spacetime

$$\nabla^{\nu}\left(R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R\right) = 0, \quad (1.18)$$

we trivially obtain:

$$\nabla^{\nu}T_{\mu\nu} = 0. \quad (1.19)$$

The last equation is the general relativistic equation for the conservation of energy momentum, in the form of a continuity equation.

³To be more mathematically precise, the relative acceleration is obtained by computing twice the directional derivative of x^{μ} along t^{μ} : $a^{\mu} = t^{\rho}\nabla_{\rho}(t^{\sigma}\nabla_{\sigma}x^{\mu})$. It turns out that for an infinitesimal displacements x^{μ} , this is proportional to $x^{\sigma}t^{\mu}t^{\nu}$, with the proportionality constant set by the Riemann tensor.

1.1.2 Linearized Gravity

General relativity is a highly non linear theory, where the combined effect on the spacetime of two distinct masses is not equal to the summed effect of the two individual masses. However, in the limit where the gravitational interaction is “weak”, the theory simplifies considerably, giving rise to a linear solution to Einstein field equation, the so-called *linearized gravity*. The weak gravity regime arises whenever the energy densities are “small” compared to the typical energy density scale in general relativity, c^4/G . In this regime, is it possible to find a coordinate system where the spacetime is nearly flat *everywhere*, i.e. $g_{\mu\nu} \simeq \eta_{\mu\nu}$, and the theory of general relativity reduces to the Newtonian theory of gravity. Far from being an unrealistic scenario, this is an excellent approximation of nature in most situations (including on Earth), as shown by the excellent agreement of the Newtonian gravity to many observations.

But the importance of linearized gravity goes well beyond the Newtonian limit of Einstein equations. Indeed, the expansion of the Einstein equation around the flat spacetime predicts the existence of gravitational waves, a purely general relativistic phenomenon. Gravitational waves are small, high-frequency perturbation to a slowly changing background metric of the spacetime. While the definition of “slowly changing background” spacetime and “high-frequency perturbation” is somewhat arbitrary, in what follow we simplify the problem and we only consider small perturbations $h_{\mu\nu}$ to a flat spacetime $\eta_{\mu\nu}$. In this limit, the spacetime metric becomes:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \quad (1.20)$$

where of course we assume $|h_{\mu\nu}| \ll 1$ everywhere, since we are only considering small perturbations. The fact that the perturbations are small implies that we can safely neglect from the equation of motion any term proportional to h^2 and $h\partial h$. Thus, as we expand the equation of motion around $\eta_{\mu\nu}$, we will only consider the first order term in $h_{\mu\nu}$ and therefore the dynamical evolution of $h_{\mu\nu}$ becomes linear with the sources.

As mentioned above, general relativity is invariant under a global coordinate transformation and we can exploit this *gauge freedom* to choose a particular reference frame in order to (i) simplify the expression of the equation of motion and (ii) highlight the physical degrees of freedom of the linear theory. Moreover, to preserve the condition $|h_{\mu\nu}| \ll 1$, any gauge coordinate transformation must be in the form of:

$$x^\mu \longrightarrow x'^\mu = x^\mu + \xi^\mu(x^\mu) \quad (1.21)$$

where the derivatives $\partial\xi^\mu$ of ξ^μ needs to be of the same order of magnitude of $h_{\mu\nu}$. Under the transformation Eq. (1.21), the metric perturbation transforms as:

$$h_{\mu\nu} \longrightarrow h'_{\mu\nu} = h_{\mu\nu} - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu) \quad (1.22)$$

where we neglect any quadratic term. Eq. (1.22) represents the gauge freedom of the linearized theory.

We can then compute the linearized versions of the Einstein equations. The computation is more conveniently done by defining

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h \quad (1.23)$$

and by choosing a gauge such that⁴

$$\partial^\nu \bar{h}_{\mu\nu} = 0. \quad (1.24)$$

With the coordinate choice above, called *Lorentz gauge*, the Einstein equations reduces to a wave equation in flat spacetime:

$$\square \bar{h}_{\mu\nu} = -16\pi T_{\mu\nu} \quad (1.25)$$

where $\square = \eta_{\mu\nu}\partial^\mu\partial^\nu$. We call gravitational wave any solution to this equation. While we will provide more insights in the next section, we can already read from the above equation that any perturbation to the spacetime (i.e. a gravitational wave) must propagate at the speed of light c .

1.1.3 Physical Degrees of Freedom of a Gravitational Wave

The general perturbation $\bar{h}_{\mu\nu}$ to the metric tensor is fully specified by ten real functions, i.e. it has ten degrees of freedom. However, as the theory must be invariant under a gauge transformation, not all those degrees of freedom are independent from each other some: we call *physical* the degrees of freedom that corresponds to a physically measurable quantity, while the rest are called *gauge* degrees of freedom. The latter are mere artifacts of the chosen coordinate system and might be set to zero with an appropriate choice of gauge. For instance the Lorentz condition Eq. (1.24) sets four constraints on the metric components, thus leaving only six independent metric components.

It turns out that a perturbation to metric propagating *in vacuum* (a gravitational wave) has only two physical degrees of freedom. This can be shown by constructing a suitable system of coordinates where $\bar{h}_{\mu\nu}$ only depends on two functions. A common choice for such gauge is the so called *transverse-traceless gauge* (TT). In this coordinate frame, in addition to the Lorentz condition, four more conditions on $\bar{h}_{\mu\nu}$ are set to give the following eight constraints:

$$\bar{h}^{0\mu} = 0; \quad \bar{h}_i^i = 0; \quad \partial^j \bar{h}_{ij} = 0 \quad (1.26)$$

⁴One can show using Eq. (1.22) that such gauge exists by showing that it is always possible to find a function ξ^μ so that the transformed $h_{\mu\nu}$ satisfy the required condition.

where we use the convention that roman indices $i, j = 1, 2, 3$ refers to the spatial components of a four-vector. In particular, note that since $\bar{h} = 0$, in the TT gauge $\bar{h}_{\mu\nu} = h_{\mu\nu}$.

In the TT gauge, a simple monochromatic plane wave, propagating in the $\hat{\mathbf{n}}$ -direction, is a trivial solution to the wave equation Eq. (1.25) and its components have the explicit expression

$$h_{\mu\nu}^{TT} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \cos(\omega(t - z/c)) \quad (1.27)$$

where we assumed, without loss of generality, that the wave is propagating along the z -axis. The quantities h_+ and h_\times are the amplitude of the so-called “plus” and “cross” polarization and they carry information about the physical degrees of freedom of a gravitational wave. To move further, we can introduce the two polarization tensors $e_{\mu\nu}^+, e_{\mu\nu}^\times$

$$e_{\mu\nu}^+ = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}; \quad e_{\mu\nu}^\times = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (1.28)$$

so that the monochromatic solution is written as a linear combination of $e_{\mu\nu}^+, e_{\mu\nu}^\times$. Moreover, we introduce the TT projector operator $\Lambda_{ij,kl}(\hat{\mathbf{n}})$, which projects into the TT gauge any plane wave solution $h_{\mu\nu}$ propagating in the $\hat{\mathbf{n}}$ -direction and expressed in the Lorentz gauge:

$$\Lambda_{ij,kl}(\hat{\mathbf{n}}) = \mathcal{P}_{ik}\mathcal{P}_{jl} - \frac{1}{2}\mathcal{P}_{ij}\mathcal{P}_{kl} \quad (1.29)$$

where we defined $\mathcal{P}_{ij} = \delta_{ij} - \hat{n}_i \hat{n}_j$ the projector on the plane orthogonal to $\hat{\mathbf{n}}$. Eq. (1.29) will be used to project any solution to the wave equation Eq. (1.25) into the convenient TT gauge.

Of course, the general solution to the wave equation is not monochromatic. However, due to the constraints introduced by the TT gauge, even the most general solution will be a linear combination of the two polarization tensors

$$h_{\mu\nu}^{TT}(t, \mathbf{x}) = h_+(t, \mathbf{x})e_{\mu\nu}^+ + h_\times(t, \mathbf{x})e_{\mu\nu}^\times \quad (1.30)$$

and the problem of modelling a gravitational wave is equivalent to the problem of determining the two polarization functions $h_+(t, \mathbf{x})$ and $h_\times(t, \mathbf{x})$. We will discuss the matter in Sec. 1.2.

1.1.4 Interaction with Test Masses

As a gravitational wave is a perturbation to the spacetime metric, which defines distances between events of the spacetime, a gravitational wave alters the distances between two free falling test masses. This alteration is detectable and forms the basis for directly detecting gravitational-wave signals, through a very precise measurement of the distance between two test masses by means of laser interferometry.

To see this effect, we note that, in the TT gauge, a test mass follows a trajectory with constant spatial coordinates x^i , since the geodesics equation trivially reduces to

$$\frac{d^2x^i}{dt^2} = 0. \quad (1.31)$$

Thus, in TT gauge, coordinates are marked by the trajectories of a set of free falling observers (test masses) and the coordinate separation between two test masses is not affected by an incoming perturbation. However, the physical distance between two test masses is affected by it. To quantify this, we define s^i to be the *proper* separation vector between two test masses, located at coordinates $\mathbf{x}_1, \mathbf{x}_2$ and separated by a constant *coordinate* separation vector $\Delta\mathbf{x} = \mathbf{x}_1 - \mathbf{x}_2$. By considering the line element ds^2 for the spacetime in the TT gauge

$$ds^2 = -dt^2 + dz^2 + (1 + h_+)dx^2 - (1 - h_+)dy^2 + (1 + h_x)dxdy \quad (1.32)$$

we obtain that at first order in h_{ij} the proper separation changes as:

$$\ddot{s}^i \simeq \frac{1}{2}\ddot{h}_{ij}^{TT} s^j \quad (1.33)$$

where \cdot denotes a time derivative of a scalar function. This shows that the effect of a gravitational is to perturb the physical distance between a pair of test masses

Even though in the TT gauge the metric perturbation has a particularly simple expression, the TT gauge is not convenient to model a physics experiment aimed at detecting gravitational waves. In a physics experiment, we use a ruler to mark our coordinates and two coordinate points ξ_1, ξ_2 are always separated by the *same* physical distance L . In this frame of reference, called *laboratory frame*, the coordinate separation $\xi^i(t)$ between two test masses is still given by Eq. (1.33)⁵

$$\ddot{\xi}^i \simeq \frac{1}{2}\ddot{h}_{ij}^{TT} \xi^j, \quad (1.34)$$

hence, unlike in the TT frame, the two test masses do change their coordinates as a consequence of a spacetime perturbation. The same expression can also be computed starting from the geodesics deviation equation Eq. (1.16) and its derivation

⁵Note that this result is non trivial. Indeed, the time derivative of the proper separation is not a gauge invariant quantity and therefore its value may change when computed in different frames.

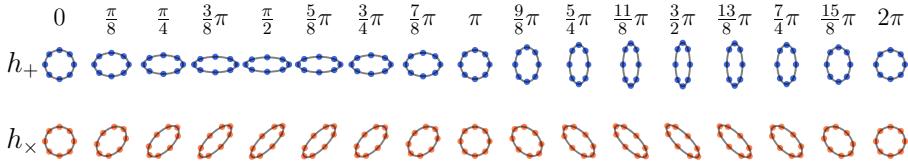


Figure 1.1: Effect on a ring of test masses of a monochromatic gravitational wave, propagating perpendicularly to the plane of the particles. For different phases ωt of the wave, we report the position of the particles, for the a plus-polarized waveform (top line) and a cross-polarized waveform (bottom line). Each particle moves according to Eq. (1.37) and Eq. (1.38) under the effect of the plus and cross polarization respectively.

assumes that the typical wavelength λ of the perturbation is much larger than the typical length L :

$$L \ll \lambda. \quad (1.35)$$

Since the unperturbed spacetime is flat and we use the distance between points to define a coordinate frame, in the laboratory frame we can apply our Newtonian intuition and the effect of a gravitational wave is equivalent to a force $\frac{1}{2}\ddot{h}_{ij}^{TT}\xi^j$ between two points of the spacetime separated by a distance ξ^i . The Newtonian force induced by an incoming gravitational wave can be measured by means of an interferometer, as explained with some details in Sec. 2.1.

We can use Eq. (1.34) to study the motion of a test particle. We will consider a plus-polarized monochromatic wave propagating on the z direction so that the x - y components of the perturbation are:

$$h_{ab}^{TT} = h_+ \sin \omega t \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (1.36)$$

where h_+ is a constant waveform amplitude. By introducing the unperturbed position ξ_0^i of a test mass and defining the perturbation $\xi^i = \xi_0^i + \delta\xi^i$, the integration of Eq. (1.34) yields:

$$\begin{aligned} \delta\xi^x(t) &= \frac{h_+}{2} \xi_0^x \sin \omega t \\ \delta\xi^y(t) &= -\frac{h_+}{2} \xi_0^y \sin \omega t \end{aligned} \quad (1.37)$$

Similarly for a cross-polarized waveform, we obtain:

$$\begin{aligned} \delta\xi^x(t) &= \frac{h_x}{2} \xi_0^y \sin \omega t \\ \delta\xi^y(t) &= \frac{h_x}{2} \xi_0^x \sin \omega t \end{aligned} \quad (1.38)$$

The effect of an incoming GW can be easily visualized by studying the motion of a set of test masses on a ring around the unperturbed position, as shown in Fig. 1.1

1.1.5 Energy of a Gravitational Wave

The problem of defining the energy carried by a gravitational wave is very subtle and highly non-trivial. In general relativity, there is no notion of gravitational energy—this fact, called *equivalence principle*, marks one of the basis for the theory. For this reason, it is impossible to associate a gauge invariant quantity to the “energy” of a gravitational field. On the other hand, we know from the previous section that a gravitational wave has the effect of moving a set of test masses and thus, being able to carry on work, it must carry some form of energy.

The solution to this problems consists in defining a somewhat arbitrary separation between a “static” background gravitational field and a dynamic gravitational-wave field. While no energy notion can be meaningfully defined for the first, for the latter we can define a gauge invariant expression for the energy. A crucial assumption is that there is a clear scale separation between the typical timescale of the background and of the gravitational-wave field. More precisely, we assume that there exist a frequency \bar{f} , which is simultaneously much higher than the typical frequency of the background f_{bg} and much smaller than the frequency of the perturbation f_{GW} : $f_{\text{bg}} \ll \bar{f} \ll f_{\text{GW}}$. If that assumption holds, we can define the energy carried by a high frequency perturbation as the average effect that the perturbation has on the background spacetime. Note that in the case of a flat background $g_{\text{bg}} = \eta$, the spatial frequency f_{bg} becomes zero, which supports our earlier assumption of classifying any metric perturbation, regardless of its frequency content, as a gravitational wave. Although we have assumed a flat background so far and will maintain this assumption throughout this work, we discuss here the more general case here for the valuable insights it offers.

To make things more precise, we can expand the Einstein equation in vacuum up around the g_{bg} solution to second order in the metric perturbation h :

$$G_{\mu\nu}(g_{\text{bg}} + h) \simeq G_{\mu\nu}^{(0)}(g_{\text{bg}}) + G_{\mu\nu}^{(1)}(h) + G_{\mu\nu}^{(2)}(h) = 0 \quad (1.39)$$

where we $G^{(i)}$ denotes the i -th term of the expansion, which depends either on the background or on the perturbation only.

The expression $G_{\mu\nu}^{(1)}(h) = 0$ is the dynamical equation for the gravitational-wave perturbation and is equivalent to Eq. (1.25) in vacuum. For this reason, $G_{\mu\nu}^{(1)}(h) = 0$ is automatically satisfied by the perturbation $h_{\mu\nu}$ and Eq. (1.39) reduces to:

$$G_{\mu\nu}^{(0)}(g_{\text{bg}}) = -G_{\mu\nu}^{(2)}(h) \quad (1.40)$$

which is equivalent to the Einstein equations for the background metric g_{bg} with a source term $-G_{\mu\nu}^{(2)}(h)$. As the tensor $G_{\mu\nu}^{(2)}(h)$ depends only on the metric perturbation and it affects the background metric, we can identify it as the energy momentum tensor of a gravitational wave. The energy momentum tensor $G_{\mu\nu}^{(2)}(h)$ is not gauge invariant, however, it becomes so after averaging over the high-frequency components: this is equivalent to integrating out degrees of freedom impossible to resolve at the large scale of observation corresponding to the background metric. We can then define the energy momentum $t_{\mu\nu}^{GW}(h)$ of a gravitational wave as:

$$t_{\mu\nu}^{GW}(h) := -\frac{1}{8\pi} \langle G_{\mu\nu}^{(2)}(h) \rangle = \frac{1}{32\pi} \langle \partial_\mu h^{\alpha\beta} \partial_\nu h_{\alpha\beta} \rangle \quad (1.41)$$

where $\langle \dots \rangle$ denotes the average of the high frequencies $f > \bar{f}$. As the expression is gauge invariant, it is also valid in the TT gauge.

Using Eq. (1.41), we can compute the energy density t_{00}^{GW} of a gravitational wave in terms of the polarizations h_+, h_\times in the TT gauge:

$$t_{00}^{GW} = \frac{1}{16\pi} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle \quad (1.42)$$

The fact that the energy density depends only on the physical degrees of freedom of a gravitational wave (and not on the gauge ones) is a further proof that a metric perturbation indeed carries energy.

1.2 Sources of Gravitational Waves

As we will show below, any accelerating object can emit gravitational waves. However, due to the smallness of the gravitational constant G entering the Einstein equations, only few accelerating systems are actually relevant for gravitational-wave astronomy. Such systems are the most “extreme” ones, where the energy densities are large enough to produce a measurable effect on the spacetime, and give rise to a large variety of signals with different morphologies.

Detectable signal and sources include bursts emitted by a core collapse supernova [87–89] or by cosmic strings produced in the early universe [90–93], long duration signals emitted by a tiny asymmetry in a fast rotating neutron star [94–97], stochastic signals [98] produced by unresolvable overlapping sources (also called astrophysical background) [99] or by perturbations in the early universe (cosmological background) [100]. In this work however, we will consider gravitational-wave signals from the coalescence of compact objects. More specifically, we will be concerned with signals produced by the coalescence of heavily precessing black hole binaries. This will be the focus of the next section. In this section, we describe the mathemati-

cal formalism needed to compute the gravitational wave emitted by a given physical system.

1.2.1 The Quadrupole Formula

To compute the gravitational waves emitted by a given source, we need to solve the wave equation Eq. (1.25) for the metric perturbation $h_{\mu\nu}$. As discussed above, we will work in the TT gauge, which is particularly convenient to describe the waveform propagation. From the structure of the wave equation, we already know that any solution $h_{ij}^{TT}(t, \mathbf{x})$ must satisfy $h_{ij}^{TT}(t, \mathbf{x}) = f_{ij}(t \pm |\mathbf{x}|)$, where the details of the functions f_{ij} depends on the source. The solution for a particular source can be readily obtained using the Green's function $G(x - x')$ of the d'Alembert operator \square :

$$G(x - x') = -\frac{1}{|\mathbf{x} - \mathbf{x}'|} \delta(t - t' - |\mathbf{x} - \mathbf{x}'|) \quad (1.43)$$

where the time $t - |\mathbf{x} - \mathbf{x}'|$ is also called retarded time. Using the fact that the Green's function satisfies

$$\square G(x - x') = \delta^{(4)}(x - x'), \quad (1.44)$$

we are able to obtain the general solution for the metric perturbation projected on the TT gauge:

$$h_{ij}^{TT}(t, \mathbf{x}) = 4\Lambda_{ij,kl}(\hat{\mathbf{n}}) \int d^3x' \frac{1}{|\mathbf{x} - \mathbf{x}'|} T_{kl}(t - |\mathbf{x} - \mathbf{x}'|, \mathbf{x}'). \quad (1.45)$$

This expression, coupled with the knowledge of the energy momentum tensor $T_{ij}(t, \mathbf{x})$ fully solves the problem of computing the gravitational-wave emission of a source.

Clearly, in most practical applications Eq. (1.45) is impossible to solve analytically, or even to tackle numerically. For this reason, we are interested on computing a first order approximation for the solution, which is valid for *slowly moving* sources, characterized by a velocity v negligible with respect to the speed of light. In this scenario, we are interested in the behaviour of the waveform far away from the source, i.e. at large $R = |\mathbf{x}|$. We begin by computing the Fourier transform [101] $\tilde{h}_{ij}^{TT}(\omega, \mathbf{x})$ of $h_{ij}^{TT}(t, \mathbf{x})$ to obtain:

$$\tilde{h}_{ij}^{TT}(\omega, \mathbf{x}) = \frac{4}{\sqrt{2\pi}} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \int dt d^3x' e^{i\omega t} \frac{1}{|\mathbf{x} - \mathbf{x}'|} T_{kl}(t - |\mathbf{x} - \mathbf{x}'|, \mathbf{x}') \quad (1.46)$$

$$= \frac{4}{\sqrt{2\pi}} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \int dt_{\text{ret}} d^3x' e^{i\omega t_{\text{ret}}} \frac{e^{i\omega|\mathbf{x}-\mathbf{x}'|}}{|\mathbf{x} - \mathbf{x}'|} T_{kl}(t_{\text{ret}}, \mathbf{x}') \quad (1.47)$$

$$= 4\Lambda_{ij,kl}(\hat{\mathbf{n}}) \frac{e^{i\omega R}}{R} \int d^3x' \tilde{T}_{kl}(\omega, \mathbf{x}') \quad (1.48)$$

where to obtain the second line we performed a change of variables $t_{\text{ret}} = t - |\mathbf{x} - \mathbf{x}'|$ and in the third we introduced the frequency domain energy momentum tensor $\tilde{T}_{kl}(\omega, \mathbf{x}')$ and we made the approximation $\frac{e^{i\omega|\mathbf{x} - \mathbf{x}'|}}{|\mathbf{x} - \mathbf{x}'|} \simeq \frac{e^{i\omega R}}{R}$. The latter approximation is a consequence of our assumption of a far-away source ($|\mathbf{x} - \mathbf{x}'| \simeq R$) and of a slowly-moving source ($\omega|\mathbf{x}'| \ll 1$) and corresponds to considering the first order term of an expansion of $|\mathbf{x} - \mathbf{x}'| \simeq R + \delta R + \dots$

To move further, we note that from the conservation of the energy momentum $\partial_\nu T^{\mu\nu} = 0$, we have the following identity in Fourier space

$$-i\omega \tilde{T}^{0\mu} = \partial^i \tilde{T}^{0i}. \quad (1.49)$$

Therefore, we can integrate $\int d^3x' \tilde{T}_{kl}(\omega, \mathbf{x}')$ by parts twice and discard boundaries terms to obtain:

$$\int d^3x' \tilde{T}^{ij}(\omega, \mathbf{x}') = -\frac{\omega^2}{2} \int d^3x' \tilde{T}^{00}(\omega, \mathbf{x}') x^i x^j. \quad (1.50)$$

The integral expression in the right hand side is equivalent to the *mass quadrupole* tensor $M^{ij}(t)$, which has the following expression in time domain:

$$M^{ij}(t) = \int d^3x' T^{00}(t, \mathbf{x}') x^i x^j. \quad (1.51)$$

By plugging Eq. (1.50) into Eq. (1.48) and computing the inverse Fourier transform, we obtain:

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{4}{R} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \ddot{M}^{kl}(t). \quad (1.52)$$

This equation is the famous quadrupole formula and it relates the gravitational-wave emission of a compact system to the second time derivative of its mass quadrupole Eq. (1.51). It is valid in the far zone only for slowly moving sources and it is not able to resolve the fine details of a generic source. While the formula can provide a satisfactory approximation to Eq. (1.45), it is only the first term of an expansion of $|\mathbf{x} - \mathbf{x}'|$ around R . In the next section, we will discuss a more systematic approach to the expansion of Eq. (1.45), leading to the so-called *multipole expansion* of gravitational-wave radiation.

1.2.2 Multipole Expansion of the Source and of $\mathbf{h}(t)$

The quadrupole formula Eq. (1.52) relies on two approximations, namely that the sources are slowly moving and that the observer is far away from the source. Under this assumptions, we approximated $|\mathbf{x} - \mathbf{x}'| \simeq R$ in Eq. (1.46) and obtained a strikingly simple expression. However, the solution we obtained is only the first term of an infinite series for h_{ij}^{TT} . Performing a systematic expansion is a complicated task,

outside the scope of this work, thus we limit ourself to summarize a few interesting results.

The expanded solution to the wave equations can be written in terms of the statistical moments, also called *multipoles*, $M^{i_1 \dots i_N}$ of the energy density distribution T^{00}

$$M^{i_1 \dots i_N} = \int d^3x T^{00}(t, \mathbf{x}) x^{i_1} \dots x^{i_N} \quad (1.53)$$

and of the multipoles $P^{i, i_1 \dots i_N}$ of the linear momenta T^{0i} , also called *current multipoles*,

$$P^{i, i_1 \dots i_N} = \int d^3x T^{0i}(t, \mathbf{x}) x^{i_1} \dots x^{i_N}. \quad (1.54)$$

As we already seen, the first term of the expansion only depends on the second derivative of the mass multipole M^{ij} . Higher order terms will involve a combination of high order time derivatives of $M^{i_1 \dots i_N}$ and $P^{i, i_1 \dots i_N}$. For instance, at second order h_{ij}^{TT} is given by:

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{4}{R} \Lambda_{ij,kl}(\hat{\mathbf{n}}) \left\{ \ddot{M}^{kl}(t) + n_m \left[\frac{1}{6} \ddot{M}^{klm} + \frac{1}{3} \left(\ddot{P}^{k,ml} + \ddot{P}^{l,mk} - 2\ddot{P}^{m,lk} \right) \right] + \dots \right\}. \quad (1.55)$$

This expression is called the *multipole expansion of the source* [102].

It turns out that it is more convenient to express the spatial dependence of the solution in spherical coordinates (r, θ, ϕ) and express the solution in terms of a basis of complex functions defined on the sphere, called spin-2 spherical harmonics $Y_{\ell m}(\theta, \phi)$ [103]. The spin-2 spherical harmonics have the property that:

$$Y_{\ell-m} = (-1)^m Y_{\ell m}^*. \quad (1.56)$$

Moreover, they define a set of multi-indices tensors $\mathcal{Y}_{i_1 \dots i_N}^{\ell m}$, called spherical tensors, as follows:

$$Y_{\ell m}(\theta, \phi) = \mathcal{Y}_{i_1 \dots i_\ell}^{\ell m} n_{i_1} \dots n_{i_\ell} \quad (1.57)$$

where the vector $\hat{\mathbf{n}}$ defines the z-axis of the spherical coordinate system chosen. Note that the indices ℓm are not *covariant* indices, i.e. they are not related to any spacetime coordinate. As a consequence, we make no distinction between upper and lower indices as we did for the covariant indices in the previous sections. The spherical tensors can be used to project a generic tensor $A^{i_1 \dots i_N}$ onto the basis of spherical harmonics by computing the scalar product

$$A^{\ell m} = A^{i_1 \dots i_N} \mathcal{Y}_{i_1 \dots i_N}^{\ell m}. \quad (1.58)$$

The quantity $A^{\ell m}$ is called the spherical component of the tensor $A^{i_1 \dots i_N}$. The decomposition is very convenient since in spherical coordinates, the general solution for h_{ij}^{TT} in the far zone can be written [85, 104, 105]:

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{1}{r} \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{m=\ell} \left[U_{\ell m} \mathbf{T}_{ij}^{E2, \ell m}(\theta, \phi) + V_{\ell m} \mathbf{T}_{ij}^{B2, \ell m}(\theta, \phi) \right] \quad (1.59)$$

where the functions $U_{\ell m}$ and $V_{\ell m}$ are directly related to the mass and linear momentum multipoles of the source:

$$U_{\ell m} = \frac{d^{\ell}}{dt^{\ell}} M^{i_1 \dots i_{\ell}} \mathcal{Y}_{i_1 \dots i_{\ell}}^{\ell m} \quad (1.60)$$

$$V_{\ell m} = \frac{d^{\ell}}{dt^{\ell}} \epsilon_{ijk} P^{j, k i_1 \dots i_{\ell-1}} \mathcal{Y}_{i_1 \dots i_{\ell-1}}^{\ell m} \quad (1.61)$$

and ϵ_{ijk} is the totally antisymmetric symbol. The tensors $\mathbf{T}_{ij}^{E2, \ell m}(\theta, \phi)$ and $\mathbf{T}_{ij}^{B2, \ell m}(\theta, \phi)$ carry the angular dependence of the solution and are defined in terms of the spherical harmonics $Y_{\ell m}(\theta, \phi)$ [85, 104]. The two gravitational-wave degrees of freedom h_+, h_{\times} can be extracted from Eq. (1.59) as

$$h_+ + i h_{\times} = h_{ij}^{TT} n^i n^j = \frac{1}{r} \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{m=\ell} h_{\ell m} Y_{\ell m}(\theta, \phi). \quad (1.62)$$

The functions $h_{\ell m}$ are called *modes* of a waveform and are a function of only $U_{\ell m}$ and $V_{\ell m}$:

$$h_{\ell m} = \frac{1}{\sqrt{2}} (U_{\ell m} + i V_{\ell m}). \quad (1.63)$$

The result above connects the multipole expansion of the source Eq. (1.55) to the expansion in spherical harmonics of the gravitational wave $h_+ + i h_{\times}$ in the far zone. For the purpose of modelling, it is customary to write each mode $h_{\ell m}$ in terms of its amplitude $A_{\ell m}$ and phase $\varphi_{\ell m}$:

$$h_{\ell m}(t) = A_{\ell m}(t) e^{i \varphi_{\ell m}(t)}. \quad (1.64)$$

Moreover, as each mode corresponds to a different order in the expansion of the source, different modes have different typical magnitudes, with $\ell = |m| = 2$ mode corresponding to the dominant quadrupole emission. While most detected binary systems reveal only the dominant mode in observations, evidence suggests that for a few asymmetric systems, higher order modes (HMs) beyond the dominant one are also detectable [106, 107].

1.3 Black Hole Binaries

A binary black hole (BBH) is a binary system composed by two black holes. As the two black holes orbit around each other, they emit gravitational radiation, which dissipates orbital energy causing the orbit to shrink until the two black holes eventually merge. The gravitational waves emitted by this process is detectable by current interferometers. After a brief discussion of black holes, black hole binaries and their formation, we dive into some details of their orbital evolution. Finally we will review the role of the two BH spins in determining the dynamics of the binary, eventually discussing the spin-induced precession motion of the orbital plane, whose detection is the main focus of this work.

1.3.1 Black Hole and Black Hole Binaries in Astrophysics

A black hole is a region of the spacetime which is not able to communicate with the rest of the space time. Indeed, inside this region the gravitational interaction is so strong that no object or light sign can escape from it. Inside a black there is a *singularity*, which is a region that the general relativity is unable to appropriately describe and for which predicts an infinite curvature. Every object inside the black hole will eventually reach the singularity, or more formally, every geodesics inside the black hole end inside the singularity. The nature of the singularity is still unknown and probably its understanding requires a new theory, possibly incorporating the theory of quantum mechanics within the framework of general relativity.

A black hole arises naturally as a solution of the Einstein equations and it can be uniquely described by its mass M and possibly by its spin vector \mathbf{S} and its electric charge q . The spin vector is bounded by the mass of the black hole and must satisfy $S \leq M^2$; it is also common to define a *dimensionless* spin vector $\mathbf{s} = \mathbf{S}/M^2 \in [0, 1]$. A chargeless and spinless black hole is described by the Schwarzschild solution [108], while a rotating chargeless black hole is well modeled by the Kerr solution [109, 110]. Finally, a charged black hole is described by the Reissner–Nordström [111] solution in the spinless case and by Kerr–Newman [112, 113] solution in the spinning case. Each of this solution assumes that the black hole is immersed in vacuum and it is isolated (technically, it assumes an asymptotically flat spacetime) and provides a metric tensor $g_{\mu\nu}$ modelling the spacetime outside and inside the black hole [114].

Black holes are considered to be one of the end point of stellar evolution. After the termination of the nuclear fusion processes supporting a star, the matter composing a star collapses under its own gravity forming a *compact object* [115]. Depending on the mass of the star and on its composition, the star will collapse into a white dwarfs (with a maximum mass [116] around $1.4 M_{\odot}$), into a neutron stars (with an

unknown maximum mass [117, 118] between $2.2 - 2.9 M_{\odot}$) or into a black hole. The black holes resulting from stellar collapse typically have masses from $\sim 3 M_{\odot}$ up to $\sim 50 M_{\odot}$ [115, 119]. Besides stellar black holes, supermassive black holes [120, 121] have also been observed at the center of some galaxies [122] with masses of hundreds of thousands solar masses. They play an active role in the dynamics of the galaxy and their origin is still under study [123–125].

As for other objects in the universe, stellar mass black holes can be found in binary systems. Such binaries are thought to have primarily two formation channels [61]:

- *Common evolution*: two stars in a binary system end their life cycle and both of them collapse into a black hole while still orbiting around each other. Black hole binaries formed in this way tend to have equal mass objects with highly aligned spins, hence showing little sign of precession [126–131].
- *Dynamical capture*: a black hole falls into the potential well of another isolated black hole. Systems formed in this way have more asymmetric masses and show larger spin mis-alignment, possibly showing a larger amount of precession [132–139].

Black hole binaries are routinely detected through their gravitational-wave emission, whose study allows us to gain important information about the population of BBH and their formation with impacts in the fields of cosmology [29] or of tests of the theory of general relativity [32–34]. In what follows we describe the orbital evolution of a black hole binary and its gravitational-wave emission.

1.3.2 Orbital Evolution of a Black Hole Binary

In the limit where the two objects are far apart, the dynamics of the system is well approximated by the Kepler's law and the emission of a binary system is well approximated by the quadrupole formula. Moreover, we assume that the orbital shrinking due to gravitational-wave energy loss, also called *radiation reaction*, happens at a slow rate as compared to the orbital frequency. With this three assumptions, we can obtain a simple approximate expression, usually called Newtonian approximation, for the time evolution of a binary black hole and for the GW radiation emitted [140]. The slow orbital shrinking, which we are modelling with the Newtonian approximation, is called *inspiral*.

We assume that the two black holes with masses m_1, m_2 follows a uniform circular trajectory with frequency $\omega_{\text{orb}}/2\pi$. As it is well known from Newtonian gravity, such

system can be mapped into a single mass with mass $\mu = \frac{m_1 m_2}{m_1 + m_2}$ with trajectory $\mathbf{x}(t)$

$$x(t) = R \cos(\omega_{\text{orb}} t + \pi/2) \quad (1.65)$$

$$y(t) = R \sin(\omega_{\text{orb}} t + \pi/2) \quad (1.66)$$

$$z(t) = 0 \quad (1.67)$$

and with an energy density $\rho(t, \mathbf{x}) = \mu \delta^{(3)}(\mathbf{x}(t) - \mathbf{x})$. This allows us to compute the quadrupole moment M^{ij} Eq. (1.51) and its second time derivative \ddot{M}^{ij} , whose only non-zero components are:

$$\ddot{M}_{11} = 2\mu R^2 \omega_{\text{orb}}^2 \cos 2\omega_{\text{orb}} t \quad (1.68)$$

$$\ddot{M}_{12} = 2\mu R^2 \omega_{\text{orb}}^2 \sin 2\omega_{\text{orb}} t. \quad (1.69)$$

Plugging this expression inside the quadrupole formula Eq. (1.52) and choosing a generic direction of propagation $\mathbf{n} = (\sin \iota \cos \phi, \sin \iota \sin \phi, \cos \iota)$, parameterized by two angles ι and ϕ , we obtain an expression for the two polarizations:

$$h_+(t; r, \iota, \phi) = \frac{1}{r} \mu R^2 \omega_{\text{orb}}^2 \frac{1 + \cos^2 \iota}{2} \cos(2\omega_{\text{orb}} t_{\text{ret}} + 2\phi) \quad (1.70)$$

$$h_\times(t; r, \iota, \phi) = \frac{1}{r} \mu R^2 \omega_{\text{orb}}^2 \cos \iota \sin(2\omega_{\text{orb}} t_{\text{ret}} + 2\phi) \quad (1.71)$$

where t_{ret} denotes retarded time. The angles ι and ϕ are also known as *inclination* angle and *reference phase*. By using the Kepler's relation between orbital frequency ω_{orb} and orbital radius R of the two objects

$$\omega_{\text{orb}}^2 = \frac{m_1 + m_2}{R^3} \quad (1.72)$$

and introducing the *chirp mass* M_c of the system:

$$M_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}} \quad (1.73)$$

we obtain

$$h_+(t) = \frac{1}{r} M_c^{5/3} (\pi f_{\text{gw}})^{2/3} \frac{1 + \cos^2 \iota}{2} \cos(2\pi f_{\text{gw}} t_{\text{ret}} + 2\phi) \quad (1.74)$$

$$h_\times(t) = \frac{1}{r} M_c^{5/3} (\pi f_{\text{gw}})^{2/3} \cos \iota \sin(2\pi f_{\text{gw}} t_{\text{ret}} + 2\phi) \quad (1.75)$$

where we defined the frequency of the gravitational wave as *twice* the orbital frequency: $2\pi f_{\text{gw}} = \omega_{\text{gw}} = 2\omega_{\text{orb}}$. The two polarizations h_+ and h_\times are identical except for an inclination dependent amplitude scaling and $\pi/2$ phase shift. As we will see in Ch. 3, this has a deep impact on the way we search for signals for black hole binaries. Besides the chirp mass Eq. (6.10), we may also introduce the total mass M of

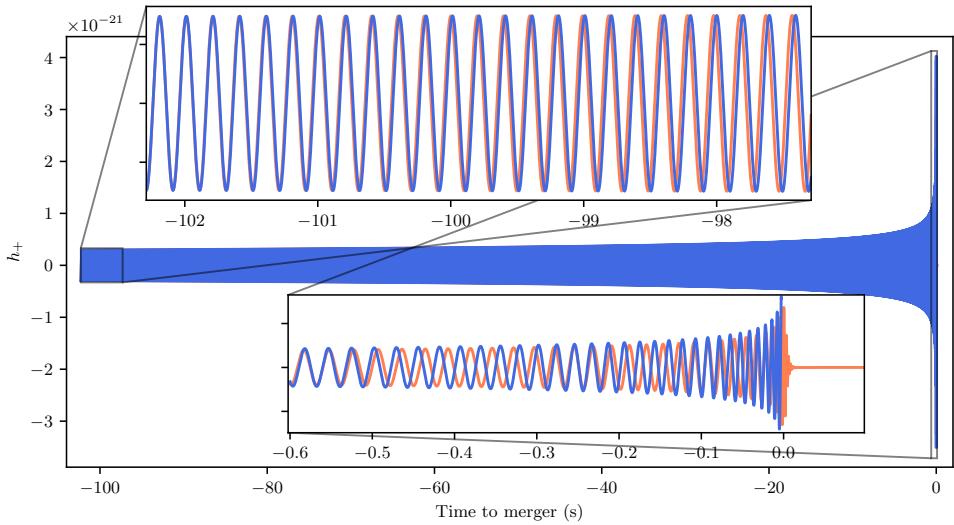


Figure 1.2: Waveform emitted by a binary system according to the Newtonian approximation Eqs. (1.85-1.86) (in blue), compared to a more realistic prediction generated with the state-of-the-art `IMRPhenomT` waveform approximant (underlaid in orange). The two waveforms refer to a system with masses $m_1 = 30 M_\odot$ and $m_2 = 10 M_\odot$, seen from a distance of 100 Mpc and with zero inclination, and generated from a minimum frequency $f_{\min} = 5$ Hz.

the system $M = m_1 + m_2$, the mass ratio $q = m_1/m_2 \geq 1$ and the symmetric mass ratio $\eta = \frac{\mu}{M} = \frac{m_1 m_2}{(m_1 + m_2)^2}$. Such quantities defines a common parametrization for the masses m_1, m_2 of the binaries, which may be useful in certain cases. Associated to the orbital rotation, the binary possesses an orbital angular momentum L :

$$L = \mu R^2 \omega_{\text{orb}} = \mu \sqrt{MR} = \mu M^{2/3} \omega_{\text{orb}}^{-1/3} \quad (1.76)$$

which plays an important role in the description of the phenomenon of precession, as we will see in the next section.

According to Eq. (1.42), a binary system will radiate energy in the form of gravitational radiation. The total radiated power P , integrated over the solid angle, can be computed from the two polarizations to be:

$$P = \int d\Omega \frac{dP}{d\Omega} = \frac{32}{5} (M_c \pi f_{\text{gw}})^{10/3}. \quad (1.77)$$

The source of radiated energy is the orbital energy

$$E_{\text{orb}} = -\frac{m_1 m_2}{2R} = -\left(\frac{\pi^2}{8} M_c^5 f_{\text{gw}}^2\right)^{1/3} \quad (1.78)$$

which decreases to compensate for the radiated energy. Therefore, from the conservation of energy

$$\dot{E}_{\text{orb}} + P = 0, \quad (1.79)$$

we can obtain the rate of orbital shrinking, or equivalently the rate of increase in the wave frequency

$$\dot{f}_{\text{gw}} = \frac{96}{5} \pi^{8/3} M_c^{5/3} f_{\text{gw}}^{11/3}. \quad (1.80)$$

This can be integrated in time to give the frequency evolution of the binary

$$f_{\text{gw}}(\tau) = \frac{1}{\pi} \left(\frac{5}{256} \frac{1}{\tau} \right)^{3/8} M_c^{-5/8} \quad (1.81)$$

where τ is the time to coalescence $\tau = t_{\text{coal}} - t$. To gain a better intuition of the magnitude of the quantities involved, we can discard the natural units and express time in seconds and the chirp mass in solar masses, resulting in the following expression:

$$f_{\text{gw}}(\tau) \simeq 151 \text{ Hz} \left(\frac{1 \text{ M}_\odot}{M_c} \right)^{5/8} \left(\frac{1 \text{ s}}{\tau} \right)^{3/8}. \quad (1.82)$$

The frequency evolution is only valid in the so called *quasi circular* regime, where

$$\dot{f}_{\text{gw}} \ll f_{\text{gw}}^2, \quad (1.83)$$

as the model for the gravitational-wave emission Eq. (1.68) is only valid for the circular case.

Under this assumption, the phase of the gravitational wave emitted by the binary can be obtained by integrating the instantaneous frequency Eq. (1.81) as a function of time to obtain the phase $\varphi(\tau)$ of the waveform

$$\varphi(\tau) = \int_{\tau_0}^{\tau} d\tau 2\pi f_{\text{gw}}(\tau) = 2 \left(\frac{\tau}{5M_c} \right)^{5/8} + \varphi_0 \quad (1.84)$$

where φ_0 is an integration constant which can be freely chosen. By plugging the expression for the phase into the waveform polarizations Eqs. (1.74-1.75), we obtain the following expression for the gravitational waves emitted by a binary system, under the Newtonian approximation:

$$h_+(t) = \frac{1}{r} M_c^{5/3} (\pi f_{\text{gw}}(\tau))^{2/3} \frac{1 + \cos^2 \iota}{2} \cos(\varphi(\tau) + 2\phi) \quad (1.85)$$

$$h_{\times}(t) = \frac{1}{r} M_c^{5/3} (\pi f_{\text{gw}}(\tau))^{2/3} \cos \iota \sin(\varphi(\tau) + 2\phi). \quad (1.86)$$

An example of such waveform is shown in Fig. 1.2 and it is compared with the same waveform computed with the state-of-the-art approximant `IMRPhenomT` [48], which

provides a more realistic waveform. Note that while the two waveforms start in phase with each other, after some cycles, the Newtonian approximation (in blue) is not able to track the more accurate phase evolution obtained by the phenomenological model (in orange). The dephasing is slow during the inspiral phase, where the Newtonian approximation retains some accuracy. During the late inspiral, the dephasing becomes much faster, indicating that the simple assumptions made for our computation are not valid for signals close to the merger of the two black holes.

It is instructive to compute the order of magnitude of the polarizations in a realistic scenario. If we consider two black holes, each with mass $10 M_{\odot}$, which are 1 s away from merger and at a realistic distance of $r = 100$ Mpc, we see that

$$h_+ \sim 10^{-21}$$

This justifies the assumption of a small perturbation to the spacetime! Moreover, it sets an ambitious sensitivity requirement for a ground based detector, such as LIGO or Virgo, designed to detect such signals (among others).

Towards a more realistic model As the two objects get closer and the orbital shrinking gets faster, the Newtonian approximation is not valid anymore and more sophisticated calculations must be performed to obtain an accurate prediction for the orbital dynamics and for the gravitational-wave emission. The computation is organized through an expansion of the Einstein equation for small velocities v (or equivalently, for weakly gravitationally bounded systems), called post-Newtonian (PN) expansion [141–143]. The PN expansion adds some correction to the amplitude and the phase of the waveform Eqs. (1.85-1.86) as a series expansion in power of v . The main physical picture however does not change: the gravitational-wave emission is the main source of orbital energy loss, which causes shrinking of the orbit.

The PN expansion breaks down in the very last few cycles of the binary. During this phase, the two black holes are very close to each other and they eventually merge into a single black hole - this is the *merger* phase of the binary evolution. In this regime, gravity is so strong that the dynamical evolution and the GW emission can be only obtained by solving the full Einstein equations with expensive numerical methods. This is one of the goals of *numerical relativity* [144–148].

The black hole produced by the coalescence of the two objects is usually in a perturbed state, where part of the orbital energy is stored into the excited modes of the event horizon. As this perturbed state is unstable, it promptly decays into the ground state (the Kerr's solution) by emitting gravitational waves. This is the *ringdown* phase. As the perturbations to the Kerr's solutions are expected to be “small”, a perturbative approach can be employed to describe the dynamics and GW emission, by expressing the perturbations as a superposition of *quasi-normal modes* [149]. The study of the

quasi-normal modes can provide invaluable information about the event horizon of a black hole, possibly leading to a better understanding of gravity in the strong regime.

An accurate prediction of the BBH gravitational-wave emission, incorporating not only the inspiral but also the merger and ringdown, is essential for most data analysis purposes. For this reason, many models for the waveform have been developed over time, in order to incorporate all the available knowledge of the dynamics of a binary black hole. Such models can be categorized in two families, both able to achieve state of the art accuracy in their predictions. The first family relies on the Effective One Body (EOB) formalism [52, 150–156], which maps the complicated general relativistic binary system into a problem governed by an effective Hamiltonian. EOB models tend to be accurate but are quite costly to generate, since for each waveform one needs to solve the Hamiltonian equation of motion. On the other hand, the phenomenological waveforms [47, 48, 157–159] rely on PN analytical expressions to model the inspiral, and on fits to numerical simulations to describe the intermediate and merger-ringdown regimes. They tend to be faster to evaluate than the EOB models. Both families, EOB and phenomenological, need to be calibrated with numerical relativity waveforms, computed by directly solving the Einstein equations in discretized form. The calibration makes sure that a model retains its accuracy even close to merger, where approximate treatments such as the post-Newtonian or EOB formalisms are no longer applicable.

Eccentricity So far we have consider binaries in circular orbits. However, in the most general solution the trajectory of a binary follows an ellipse: the system is then said to have an *eccentric* orbit. In this case, the properties of the orbit are not only characterized by the angular velocity ω but also by the eccentricity e of the system and the mean periastron anomaly a . As the latter two are time-varying quantities, they must usually be specified at a given reference orbital frequency. The eccentricity has the effect of increasing the power emitted by the binary Eq. (1.77), hence speeding up the binary evolution.

A closer study of the orbital dynamics in the eccentric case reveals that, already at the first order, the eccentricity of a binary always tends to *decrease* with time, i.e. $\frac{de}{dt} < 0$. In other words, the effect of the back-reaction of the system to a GW is to circularize the orbit. This process is so efficient that the orbit circularization usually happens well before the merger, or equivalently, the time-scale of the circularization is smaller than the merger timescale. This means than most of the BBH observed by ground based detectors will be in circular orbits. This justifies the fact that in most of the GW data analysis for compact binary coalescences, the effect of eccentricity is neglected and only systems on a circular orbits are considered. Recent work has been

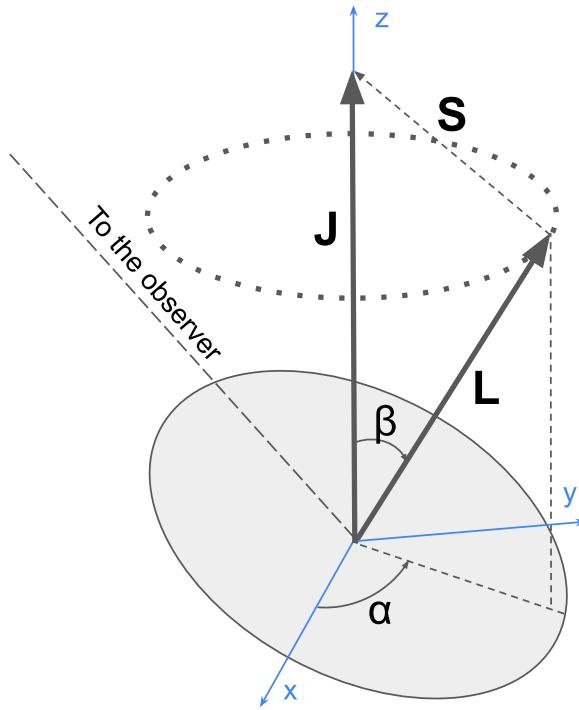


Figure 1.3: Schematic representation of the orbital plane of a precessing binary system with the orbital angular momentum \mathbf{L} , the total spin vector \mathbf{S} and the total angular momentum \mathbf{J} . We also report the angles α and β , which parameterize the precession of the orbital angular momentum \mathbf{L} around the \mathbf{J} -axis with respect to an inertial Cartesian frame.

focused on accurately modelling the effect of the eccentricity [154, 160, 161] and on incorporating its effect in the data analysis [162–166].

1.3.3 Precession in Black Hole Binaries

One interesting effect arising from the Einstein equations is the dependence of the dynamics (hence of the waveform) on the two black hole spins $\mathbf{S}_1, \mathbf{S}_2$. A simple way to understand this is to consider that, in general, the total angular momentum vector $\mathbf{J} = \mathbf{L} + \mathbf{S}_1 + \mathbf{S}_2$ retains an approximately *constant direction*, while decreasing in magnitude due to the decrease in $|\mathbf{L}|$, which in turn is caused by the orbital shrinking. If the two spins are aligned with each other and with the direction of the orbital angular momentum $\hat{\mathbf{L}}$, the conservation of the direction of \mathbf{J} implies that the orbital plane as well as the two spin vectors remain constant throughout the inspiral. Moreover, the relative orientation of the two spin vectors affects the dynamics of the system.

Heuristically, a spin vector anti-parallel to $\hat{\mathbf{L}}$ enhances the GW emission, hence it accelerates the orbital shrinking. On the contrary, two spins parallel with $\hat{\mathbf{L}}$ slow down the GW emission and the orbital shrinking. In this aligned-spin case, an interesting quantity is the *effective spin parameter* χ_{eff}

$$\chi_{\text{eff}} = \frac{m_1 s_{1z} + m_2 s_{2z}}{m_1 + m_2}, \quad (1.87)$$

which encodes the dominant spin dependence of the waveform. Consequently, the effective spin parameter is typically measured more accurately in GW observations than the individual spins.

A more interesting scenario involves the case where the two spins are mis-aligned with each other and/or with the orbital angular momentum, as first discussed in [39]. In that case, due to conservation of the direction $\hat{\mathbf{J}}$ of the orbital angular momentum, the orbital angular momentum \mathbf{L} and the two spin vectors interact with each other non-trivially and the binary experiences a gradual rotation of the orbital plane (whose normal direction is parameterized by $\hat{\mathbf{L}}$) around the total angular momentum $\hat{\mathbf{J}}$. This effect is called *precession* and it causes observable variations in the gravitational-wave signal emitted by a binary system, introducing characteristic modulation in amplitude and phase over time. A schematic representation of the physical system is reported in Fig. 1.3.

The strength of precession is characterized by the tilt angle β of the binary's orbit, i.e. the angle between \mathbf{J} and \mathbf{L} , and by the angular velocity Ω_p at which \mathbf{L} rotates around $\hat{\mathbf{J}}$. By calling S_{\parallel} and S_{\perp} the parallel and perpendicular component of the total spin \mathbf{S} with respect to the orbital angular momentum, the tilt angle β is approximately given by:

$$\tan \beta \simeq \frac{S_{\perp}}{\mu \sqrt{MR} + S_{\parallel}}, \quad (1.88)$$

where we used Eq. (1.76) to express the orbital angular momentum. As the orbital separation R decreases, we see that the precession opening angle *increases* throughout the orbital evolution, hence also the effect of precession on the waveforms is more visible towards the end of the inspiral. At leading order (see below for more details) the precession frequency is given by

$$\Omega_p \simeq \left(3 + \frac{3}{2q} \right) \frac{J}{R^3}, \quad (1.89)$$

where again we see that Ω_p increases for small separations, hence increasing the effect of precession as the orbital evolution goes on. We note that the physical picture described here holds only *approximately*. Indeed, higher order PN corrections causes the total angular momentum to also precess around a constant axis, introducing a *nutation* motion and yielding a more complicated phenomenology.

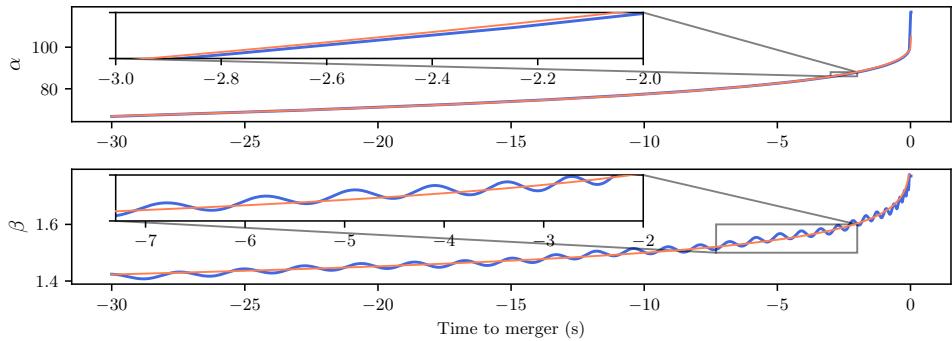


Figure 1.4: The Euler angles α and β obtained by numerically integrating Eqs. (1.96-1.97) (in blue) and with their approximate expressions Eq. (1.89) and Eq. (1.88). The fast oscillations around the approximate expression for β are caused by the nutation of \mathbf{L} along its precession motion. The angles refer to a system with masses $m_1 = 15 M_\odot$ and $m_2 = 1.5 M_\odot$ and with dimensionless spins $s_1 = (0.8, 0.0, -0.5)$ and $s_2 = (-0.7, 0.1, 0.3)$.

The situation described above is called *simple precession* and it is characterized by a non-zero total angular momentum, which provides a well defined direction for the precession axis. In the case where $\mathbf{L} \simeq -\mathbf{S}$, the system does not have an axis to rotate around and the orbital orientation \mathbf{L} rapidly changes until the system settles into a new configuration of simple precession. This phenomenon is known as *transitional precession* and, as it is expected to be rare in LIGO BBHs, in what follows we will focus only on simple precession.

Modelling the GW emission We now turn our attention to a more precise description of the gravitational waves emitted by a precessing binary. To begin, we recognize that a precessing binary system has three timescales of relevance [167–170]: the orbital timescale $t_{\text{orb}} = \omega_{\text{orb}}^{-1}$, associated to the orbital motion, the precessing timescale $t_{\text{prec}} = \Omega_p^{-1}$, associated to the change in the direction of orbital plane, and the radiation reaction timescale $t_{\text{RR}} = \omega_{\text{orb}}/\dot{\omega}_{\text{orb}}$, induced by the shrinking of the orbit due to GW energy loss. The timescale of precession t_{prec} is typically much larger than the orbital timescale t_{orb} . Moreover, the precession motion happens at a much faster rate than the shrinking of the orbit due to GW emission (radiation reaction), with timescale t_{RR} :

$$t_{\text{orb}} \ll t_{\text{prec}} \ll t_{\text{RR}}. \quad (1.90)$$

Thanks to this hierarchy of timescales, we can decouple the precessing motion from the binary evolution and from the effect of radiation reaction.

To move on, we consider a timescale \bar{t} which is short enough to neglect the pre-

cessing motion but long enough to observe many orbital cycles: $t_{\text{orb}} \ll \bar{t} \lesssim t_{\text{prec}}$. Over such timescale \bar{t} the change in the direction of the orbital angular momentum is negligible and the binary dynamics and GW emission are well approximated by those of a non-precessing system. As the orbital plane orientation changes over a timescale of t_{prec} , the system's dynamics can still be approximated by that of the same non-precessing system. However, this approximation must incorporate a rotation to account for the changing orientation. Therefore, the effect of precession can be modeled by a time dependent rotation of a suitable non-precessing binary system.

More formally, there exists a *non-inertial* reference frame where the orientation of the precessing binary *does not* change and the effect of precession on the dynamics is negligible. In such frame, called “quadrupole aligned” [171], the system dynamics and the GW emission are roughly equivalent to those of a non-precessing system with the same masses and spin magnitudes [172]. The precessing waveform can be obtained by performing a time-dependent transformation from the quadrupole aligned frame to a chosen inertial frame. The latter is typically chosen to be the frame where $\mathbf{L}(t = 0)$ is along the z -axis, also known as L_0 frame. The quadrupole aligned frame approximately coincides with frame where the total angular momentum \mathbf{J} is aligned along the z -axis and for this reason, it is sometimes referred to as J -frame.

As the quadrupole aligned frame is built to track the time evolution of the direction of the binary's orbital angular momentum $\hat{\mathbf{L}}$, it is very important to accurately model the time evolution of $\hat{\mathbf{L}}$. This is usually done within the PN framework [39,41], where it is possible to write down a set of coupled equations for the combined evolution of the orbital angular momentum \mathbf{L} and of the direction of the two spins. The equations are valid at the precessing timescale t_{prec} , as they average over many orbit cycles exploiting the hierarchy between the timescales of the problem. At 2PN order they consist in

$$\dot{\mathbf{S}}_1 = \boldsymbol{\Omega}_1 \times \mathbf{S}_1 \quad (1.91)$$

$$\dot{\mathbf{S}}_2 = \boldsymbol{\Omega}_2 \times \mathbf{S}_2 \quad (1.92)$$

$$\dot{\hat{\mathbf{L}}} = -\frac{v}{\eta M^2}(\dot{\mathbf{S}}_1 + \dot{\mathbf{S}}_2) \quad (1.93)$$

and

$$\boldsymbol{\Omega}_1 = \eta v^5 \left(2 + \frac{3}{2q}\right) \hat{\mathbf{L}} + \frac{v^6}{2M^2} \left[\mathbf{S}_2 - 3(\hat{\mathbf{L}} \cdot \mathbf{S}_2) \hat{\mathbf{L}} - \frac{3}{q}(\hat{\mathbf{L}} \cdot \mathbf{S}_1) \hat{\mathbf{L}} \right] \quad (1.94)$$

$$\boldsymbol{\Omega}_2 = \eta v^5 \left(2 + \frac{3q}{2}\right) \hat{\mathbf{L}} + \frac{v^6}{2M^2} \left[\mathbf{S}_1 - 3(\hat{\mathbf{L}} \cdot \mathbf{S}_1) \hat{\mathbf{L}} - 3q(\hat{\mathbf{L}} \cdot \mathbf{S}_2) \hat{\mathbf{L}} \right] \quad (1.95)$$

where v is the orbital velocity $v = \sqrt{M/R}$ and its time evolution can be also computed with the PN expansion [40, 173–178]. Note that if the two spins are aligned with or-

bital angular momentum, the precession frequencies Ω_i are zero. Therefore as noted before, only the *in-plane* spin components $\mathbf{S}_\perp = \mathbf{S} \times \hat{\mathbf{L}}$ contribute to precession.

Once the time evolution $\hat{\mathbf{L}}(t)$ of the orbital plane is known, it is possible to write down the explicit transformation from the quadrupole aligned frame to the L_0 frame. This is done by defining the three time dependent Euler angles $\alpha(t)$, $\beta(t)$, $\gamma(t)$, parametrizing the rotation of a rigid body [179]:

$$\alpha(t) = \arctan \frac{\hat{\mathbf{L}}_y(t)}{\hat{\mathbf{L}}_x(t)} \quad (1.96)$$

$$\beta(t) = \arccos \hat{\mathbf{L}}_z(t) \quad (1.97)$$

$$\gamma(t) = - \int_0^t dt' \dot{\alpha}(t') \cos \beta(t'). \quad (1.98)$$

Note that the β angle corresponds to the tilt angle introduced before, while the time derivative of α amounts to the precession frequency Ω_p . Approximate first order expressions for β and Ω_p respectively were previously given in Eq. (1.88) and Eq. (1.89) respectively. In Fig. 1.4, we plot the two angles α and β as computed by integrating Eqs. (1.96-1.97), together with their first order approximations.

Once the three angles are known, we can rotate the waveform in the quadrupole aligned frame (QA) to the inertial L_0 frame. The waveform rotation acts on each mode as follows:

$$h_{\ell m} = \sum_{m=-\ell}^{m=\ell} h_{\ell m'}^{QA} D_{mm'}^\ell(\alpha(t), \beta(t), \gamma(t)) \quad (1.99)$$

where $D_{mm'}^\ell(\alpha, \beta, \gamma)$ are the Wigner D matrices

$$D_{mm'}^\ell(\alpha, \beta, \gamma) = e^{im\alpha} e^{im'\gamma} d_{mm'}^\ell(\beta). \quad (1.100)$$

In this context, $d_{mm'}^\ell(\beta)$ are the real valued Wigner d matrices, which are polynomial functions of $\cos(\beta/2)$ and $\sin(\beta/2)$ and their expression can be found in [47]. The procedure outlined above goes usually under the name of *spin twist* [171, 172, 180] and is employed in most of the state-of-the-art waveform approximants [47–50, 52].

An aligned-spin system (or, approximately, a precessing system observed from the quadrupole aligned frame) exhibits a symmetry over reflection across the orbital plane, characterized by its normal vector $\hat{\mathbf{L}}$. It can be shown that this symmetry, also called *equatorial symmetry*, translates into a symmetry between different modes [181]:

$$h_{\ell m} = (-1)^\ell h_{\ell -m}^*. \quad (1.101)$$

Clearly a precessing system does not possess any equatorial symmetry and thus the symmetry between positive and negative modes is broken, as can also be seen in

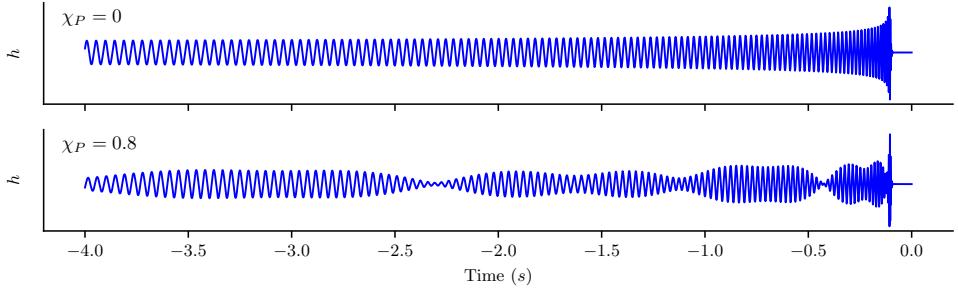


Figure 1.5: Effect of precession on a BBH signal. On the top panel we show the GW signal from an aligned-spin BBH, while in the bottom panel we show the effect of precession on the same system. On the y -axis we report the observed strain (with arbitrary units). The signal considered is characterized by $m_1 = 40 M_\odot$ and $m_2 = 3 M_\odot$ and $s_{1z} = -0.2$ and $s_{2z} = 0.7$. The precessing signal is chosen with $\chi_P = 0.8$, with precession placed only on the first BH spin. The system is in an “edge on” configuration, with inclination $\iota = \pi/2$.

Eq. (1.99). Thanks to the mode symmetry Eq. (1.101), an aligned-spin system, where only the $\ell = |m| = 2$ dominant modes are considered, is always characterized by a symmetry between the two polarizations. In frequency domain, it takes the form:

$$\tilde{h}_+ \propto i \tilde{h}_\times. \quad (1.102)$$

This expression is not valid for precessing systems and/or for systems where the higher order modes are considered. This remark has a large impact on the search for black hole binaries, as most of the searches rely on Eq. (1.102) to simplify the problem. To search for precessing signals, we need to move away from this assumption and, detailed in the next chapters, this is an important part of the original work presented in this dissertation.

Impact of precession for different systems Each precessing mode is the superposition of different aligned-spins modes $h_{\ell m}^{QA}$ and for this reason, it exhibits a characteristic modulation in the amplitude and phase, as the result of a non-trivial interference pattern. Physically, this is understood as follows: as the GW emission attains its maximum along the direction of \mathbf{L} and given that the $\hat{\mathbf{L}}$ changes with time, an inertial observers will see a time-varying amplitude as a result. Moreover, the rotation of \mathbf{L} around a constant axis will introduce a time dependent phase shift of $\alpha(t) = \int dt \Omega_p(t)$, as can also be seen in Eq. (1.99) where α enters as a phase $e^{i\alpha}$.

From a qualitative point of view, the effects of precession are more pronounced for asymmetric systems (large mass ratios) observed with an edge-on inclination (i.e. binaries observed with inclination angle $\iota \simeq \pi/2$) [73]. Clearly, larger values

of spins induce a larger amount of precession. Furthermore, the effects of precession are more pronounced in longer waveforms, primarily because there are more cycles during which precessing effects can accumulate. In typical data analysis applications, where the starting frequency of the signal is usually fixed, longer waveforms correspond to lower total masses. Consequently, with current instruments, the detection of precession is more likely in low-mass, highly asymmetric systems characterized by very large spins [74, 182].

The impact of spin misalignment can be amplified or suppressed by the inclination of the binary system, measured by the angle ι between the inclination of the observer's line of sight and the orbital angular momentum at a specific reference time. When ι is close to 0 or π , there is minimal variation in the flux of emitted gravitational waves, resulting in a limited effect of precession on the waveform. Conversely, for $\iota \simeq \pi/2$, significant variations in the GW emission occur over time. Therefore, the inclination angle is a crucial parameter to consider when searching for precessing signals.

The effect of precession is usually quantified by the *effective precession spin parameter* χ_P [172, 183], where large values of χ_P denotes large amounts of precession. Given a system with *dimensionless* spins \mathbf{s}_1 and \mathbf{s}_2 , we define $s_{1\perp} = |\mathbf{s}_1 \times \hat{\mathbf{L}}|$ as the magnitude of the in plane component of \mathbf{s}_1 and similarly $s_{2\perp}$. The effective precession spin parameter is equal to:

$$\chi_P = \max \left\{ s_{1\perp}, \frac{1}{q} \frac{4 + 3q}{3 + 4q} s_{2\perp} \right\}. \quad (1.103)$$

χ_P defines a mapping between the original binary systems (with 6 dimensional spins) and a system of reduced dimensionality, where $\mathbf{s}_1 = (\chi_P, 0, s_{1z})$ and $\mathbf{s}_2 = (0, 0, s_{2z})$. Two such systems will show similar precessing effects, since by construction the average value of the precessing frequencies Ω_i in Eqs. (1.94–1.95) are roughly the same. Despite the introduction of alternative precession measures in the literature [74, 184–187], χ_P retains its status as the most interpretable and widely utilized.

In Fig. 1.5, we show how precession changes the GW signal generated by a BBH. The upper panel reports an aligned-spin BBH, while the bottom panel shows the same signal with a non-zero χ_P . We can see how precession changes the amplitude and phase evolution of the waveform, happening with a typical timescale of t_{prec} much larger than the orbital period. This work is concerned with signals such as the one depicted in the bottom panel of Fig. 1.5.

CHAPTER 2

Detecting Gravitational Waves with Data Analysis

Some work is always required to move from theory to practice. This chapter describes how the theoretical principles introduced before come into play for the construction of very sensitive instruments for GW detection and how the instrument data can be used to extract new knowledge about the universe, through a careful data analysis. After describing the working principles of a gravitational-wave interferometer, we delve into the foundations of the techniques used to analyse interferometer data. Finally, we will give a broad overview of the variety of the scientific questions of gravitational-wave astronomy and of the data analysis techniques required to answer to them.

2.1 Interferometers for Gravitational-wave Detection

As we saw in Sec. 1.1.4, the effect of a gravitational wave is to alter the proper distance between pairs of test masses. This variation in distance can be measured by dedicated instruments, called Michelson *interferometers*. In an interferometer, a monochromatic continuous laser beam is sent to a beam splitter, which sends the beam into two orthogonal arms. At the end of each arm, a fully reflective mirror is positioned to reflect the light back to the beam splitter. The two light beams recombine into a photodetector and interfere with each other. The energy measured by the photodetector is a measure of accumulated phase difference between the optical path in the two beams

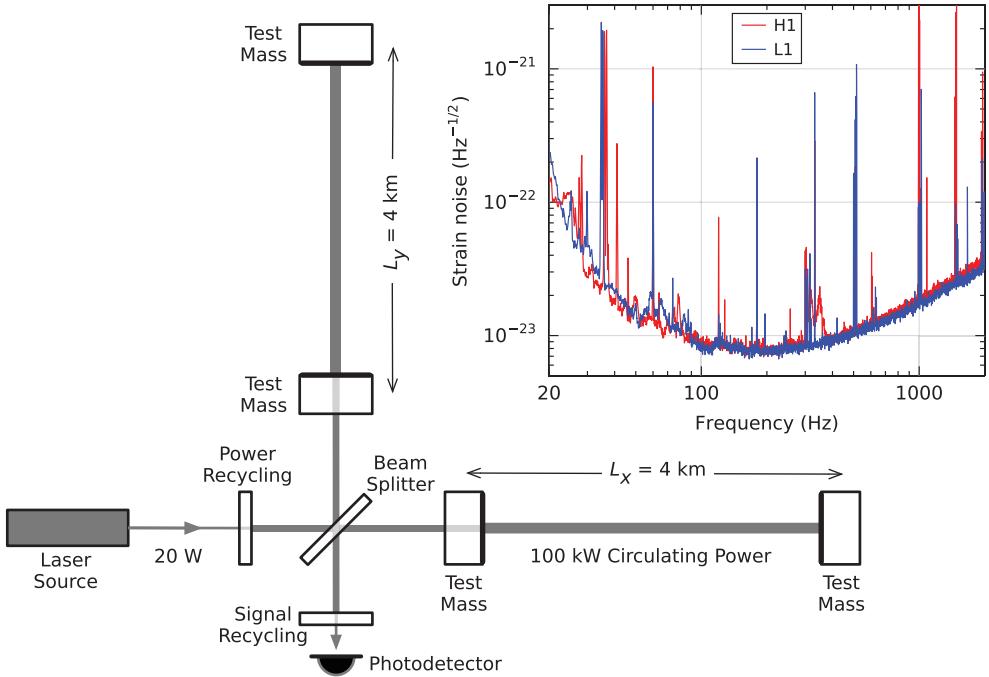


Figure 2.1: Basic setup for a Michelson interferometer to search for gravitational waves, such as LIGO or Virgo. The photodetector will measure any variation in the laser power caused by an incoming gravitational wave. In the inset, we report the typical magnitude of the instrumental noise as a function of different frequency bins, measured around the first detection of a gravitational wave in 2015. This quantity is called amplitude spectral density of the detector's output and it will be introduced in Sec. 2.2.2. Figure adapted from [188].

and it is a direct measure of the difference in distance between the two arms. As an incoming gravitational wave changes the differential arm length, the photodetector will measure a change in the collected power, possibly detecting an astrophysical signal. In this setup, the two mirrors at the end of the arms, together with the beam splitter, constitutes a set of test masses, which are free to move on the plane of the interferometer.

To make our description more quantitative, we consider an interferometer with two arms placed along the x and y -axis and with lengths at rest L_x and L_y respectively. The basic set up is described in Fig. 2.1. In what follows, we assume the wavelength λ_{GW} of the gravitational wave is much larger than the length of the detector ($L \ll \lambda_{\text{GW}}$) so that the polarization is roughly constant along the arms, or equivalently that the gravitational-wave perturbation varies slowly as compared to time of travel of the laser in the two arms ($\omega_{\text{GW}} \ll c/L$). Thanks to this assumption, the laser

can be described within a Newtonian framework (i.e. without solving the geodesics equation for the light) and we can model the experimental apparatus in the *laboratory frame*, where the effect of a gravitational wave is equivalent to a Newtonian force. For a more general solution [85], it is convenient to choose the TT gauge and to solve Eq. (1.33) for the dynamics of a test masses together with the geodesics equation for the behaviour of the laser.

After traveling back and forth into the x and y arms, two laser pulses $E_{1/2} = E_0 e^{i\omega_L t}$ of frequency ω_L observed at time t are described by:

$$E_{1/2} = \mp \frac{1}{2} E_0 e^{-i\omega_L (t - 2L_{x/y})} \quad (2.1)$$

where $t - 2L_{x/y}$ is the time where the beam entered the x/y arm and the opposite signs in the amplitude are introduced by the behaviour of the beam splitter. At any given time t , the photodetector measures the power $|E_{\text{detector}}|^2$ of the superposition $E_{\text{detector}} = E_1 + E_2$ of the two pulses:

$$|E_{\text{detector}}|^2 = E_0^2 \sin^2 [\omega_L (L_x - L_y)] \quad (2.2)$$

hence the power measured is directly related to differential arm length $(L_x - L_y)$.

By temporarily choosing the TT gauge¹ and considering the spacetime metric Eq (1.32), we can compute the length of the two arms as a function of time:

$$L_{x/y}(t) = \int_0^{L_{x/y}} \sqrt{1 \pm h_+(t)} dx \simeq \left(1 \pm \frac{1}{2} h_+(t)\right) L_{x/y} \quad (2.3)$$

where $L_{x/y}$ refers to the length of the unperturbed arms. We can then obtain a simple relation between the gravitational wave and the variation $\Delta(L_x - L_y)$ in the differential arm length:

$$\frac{\Delta(L_x - L_y)(t)}{L} = h_+(t) \quad (2.4)$$

where $L = \frac{L_x + L_y}{2}$. Note that, thanks to our assumption $L \ll \lambda_{\text{GW}}$, the polarization is constant along the arms: this allowed to simplify the integral in Eq. (2.3).

The expression above is only valid in the special case of a plus polarized wave propagating along the z direction. For a waveform with generic polarization and direction of propagation, a careful computation [14] reveals that for a Michelson interferometer the differential arm length ΔL is related to the two polarizations by a similar relation

$$\frac{\Delta(L_x - L_y)(t)}{L} = h(t) \quad (2.5)$$

¹This choice is particularly convenient as the coordinate distance between two test masses is constant with time, hence the integral has a particularly simple form. Clearly, as we are computing a scalar quantity, the result is gauge-independent.

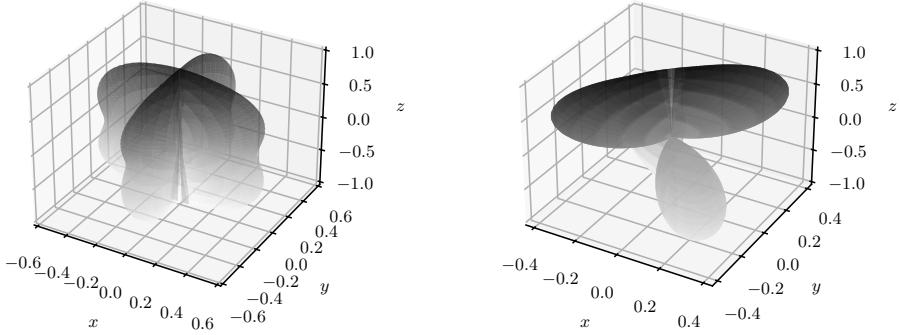


Figure 2.2: Visual representation of the antenna pattern functions F_+ , F_x for a ground based interferometer, Eq. (2.7) and (2.8) respectively. In the left panel, along each direction originating from the center and parameterized by the angles ϑ and φ , the gray-shaded surface extends to a distance $F_+(\vartheta, \varphi, 0)$ from the origin. In the right panel we make the same plot for $F_x(\vartheta, \varphi, 0)$. Note that in both cases, we set the polarization angle Ψ to zero.

where $h(t)$ is called *gravitational-wave strain* and amounts to a linear combination of the two polarizations, with the combination coefficients depending on the geometry of the instrument and the direction of propagation of the waves:

$$h(t) = F_+(\vartheta, \varphi, \Psi)h_+(t) + F_x(\vartheta, \varphi, \Psi)h_x(t). \quad (2.6)$$

In this context, ϑ and φ are the polar and azimuthal angles of the direction of propagation of the waves with respect to the normal to the detector plane, and Ψ is the *polarization angle*, characterizing an arbitrary rotation of the polarization tensors $e_{\mu\nu}^+, e_{\mu\nu}^x$ Eq. (1.28) around the direction of propagation of the wave. The functions F_+ , F_x are called *antenna pattern* functions and have the following expression:

$$F_+(\vartheta, \varphi, \Psi) = \frac{1}{2} (1 + \cos^2 \vartheta) \cos 2\varphi \cos 2\Psi - \cos \vartheta \sin 2\varphi \sin 2\Psi \quad (2.7)$$

$$F_x(\vartheta, \varphi, \Psi) = \frac{1}{2} (1 + \cos^2 \vartheta) \cos 2\varphi \sin 2\Psi + \cos \vartheta \sin 2\varphi \cos 2\Psi. \quad (2.8)$$

In Fig. 2.2, we visualize the behavior of F_+ , F_x for different directions in the sky. The antenna pattern functions defines the sensitivity of the detector towards different source locations and different gravitational-wave detectors, such as the resonant bar detectors [189], are described by different antenna pattern functions.

From Eq. (2.5) and Eq. (2.2), we see that the power measured by the photodetector depends on the gravitational-wave perturbation: this simple remark lies at the foundation of the gravitational-wave interferometry. By combining the two expres-

sions, we obtain:

$$|E_{\text{detector}}|^2 = E_0^2 \sin^2 [\omega_L(L_x - L_y) + \omega_L h(t)L]. \quad (2.9)$$

An effective choice of experimental setup will tune the arm lengths L_x, L_y in order to maximise the sensitivity to a gravitational wave.

Naively, one may think it is desiderable to maximise the change in measured power $|E_{\text{detector}}|^2$, hence setting $\omega_L(L_x - L_y) = \pi/4$ to maximise the derivative of \sin^2 . However, due to the large laser power readout, any gravitational-wave signal will be subdominant as compared to the Poisson uncertainty \sqrt{N} for the photon count N . For this reason, it is convenient to set $\omega_L(L_x - L_y) = 0$, so that the unperturbed instrument is in a state of *destructive interference*. In this case, the presence of a gravitational wave will increase the photon count well above the Poisson uncertainties. Any realistic interferometer works under this condition of destructive interference, called *dark fringe*.

2.1.1 Features of a Real Interferometer

A real ground-based interferometer, such as LIGO and Virgo, has very strict sensitivity requirements, due to the smallness of the effects it is designed to measure. As seen above, a typical gravitational-wave signal has a magnitude $h \sim 10^{-21}$, which for a interferometer with an arm of $L \simeq 4$ km translates into a change in the differential arm length of $\Delta(L_x - L_y) \sim 10^{-18}$ m. Such high precision cannot be achieved by the simple design outlined above but only with several sophisticated experimental devices.

To obtain a sensitivity improvement of several orders of magnitude, the optical path of the laser can be increased by placing a semi-reflective parallel mirror at the beginning of each arm. This configuration creates in each arm a so-called Fabry–Pérot cavity [190], where the laser bounces back and forth hundreds of times, producing a complex interference pattern. As the optical path length increases to hundreds of kilometers, effectively increasing the arm's length, the same sensitivity can be achieved with a smaller precision in the measure of $\Delta(L_x - L_y)$. As the cavity acts as a filter on the laser spectrum, the length of the cavity must be carefully chosen in order to achieve *constructive interference* at the laser operating frequency: such cavity is also called resonant. With careful computations [85], we can see that the cavity is resonant only if its length L_{cavity} satisfy:

$$\omega_L L_{\text{cavity}} \simeq n\pi \quad (2.10)$$

for any integer n .

In general such situation is not achieved, as the normal oscillation of the mirrors will prevent the experimenters to impose a distance of their choice between the mirrors of the cavity. For this reason, the movement of each mirror is controlled by magnetic attenuators, which are able to impose a small force upon them. By properly controlling all the attenuators simultaneously, it is possible to precisely control the length of the cavity and achieve constructive interference for the laser circulating inside the interferometer. This is called the *lock* state of the interferometer and it is possible thanks to an advanced control system, which quickly counterbalance any random fluctuation of the mirrors of the system. The interferometer reaches enough sensitivity to detect a gravitational wave only when it set to the lock state.

An interferometer in lock state produces a measurement $L_{\text{meas}}^{\text{diff}}(t)$ of the differential arm length $L_x - L_y$. However, the measure is affected by the mechanical correction to the free movement of the mirrors and this has to be taken into account to observe the effect of a gravitational wave $L_{\text{free}}^{\text{diff}}(t)$ on the arm length. By assuming a linear response of the detector to the attenuation, the two quantities are related in frequency domain by a simple relation:

$$\tilde{L}_{\text{meas}}^{\text{diff}}(f) = \frac{\tilde{L}_{\text{free}}^{\text{diff}}(f)}{1 + G(f)} \quad (2.11)$$

where $G(f)$ is the interferometer response function, which relates the variation of the readout to the various input forces exerted by the attenuators. The process of estimating $G(f)$ is crucial to measure a gravitational wave and it is called *calibration* of the interferometer [191–194]. We can think of the calibration process as a way to “filter out” the effect of the attenuators on the differential arm length, in order to provide a sensible estimate of a gravitational-wave signal.

Besides the Fabry–Pérot cavity and the control loop, a state-of-the-art interferometer is composed of many other devices, such as a power recycling cavity, suspension mechanisms, vacuum chambers for the laser propagation, laser squeezing or advanced coatings for the mirrors. A complete description of all of them is outside the scope of this work and the interested reader is encouraged to read [195] for more details.

2.1.2 Sources of Noise

After calibration, the instrument produces a noisy estimation $s(t)$ of the gravitational-wave strain

$$s(t) = n(t) + h(t) \quad (2.12)$$

where $h(t)$ is given by Eq. (2.6) in the case of a spacetime perturbation but of course, it might also be zero in the case where no gravitational signal is recorded. The noise

$n(t)$ is commonly assumed to be stationary over short timescales (see below for more discussion on the characterization of the noise) and it can be traced back to different sources [196, 197]:

- *Seismic noise*, arising from the displacement of the mirrors due to ground movement, produced by both natural and artificial sources. The seismic noise dominates for frequencies below 20 Hz and becomes a major limitation to the sensitivity for frequencies below ~ 5 Hz. To mitigate the effect of seismic noise, each mirror is suspended with a complicated set of pendula, which isolates the mirror from the ground and dampen the induced oscillations [198, 199].
- *Newtonian noise*, arising from the Newtonian gravitational interaction of the test masses with density waves surrounding the instrument. It mostly affects the noise budget below 1 Hz. While Newtonian noise is not currently the dominant source of noise at low frequencies, it will become so with an improved suspension mechanism and it is likely to be the dominant source of low-frequency noise in the next generation of detectors [200, 201], such as the Einstein Telescope [35, 36]. Unlike seismic noise, Newtonian noise cannot be eliminated with a suitable experimental apparatus, due to the impossibility of screening the gravitational force. Therefore, the effect of Newtonian noise has to be mitigated by properly modelling its effect on the experimental readout and by algorithmically removing it [202, 203].
- *Thermal noise*, arising from the thermally induced Brownian motion of the mirror surface. This source of noise is dominant between a few tens and a few hundreds of Hz. While the thermal noise can be mitigated by a careful choice of the mirror material, a great improvement comes from putting the mirror at cryogenic temperatures. This is the strategy adopted in the KAGRA detector [13, 198] and in the proposed Einstein Telescope.
- *Shot noise*, due to the error in the photon counting by the photodetector [204]. Due to the discrete nature of photons, the photon hits the detector following a Poisson stochastic process. The uncertainties in the photon count, translates into uncertainties in the measured power and dominate the noise budget for frequencies above 100 Hz. Since the shot noise originates from a fundamental quantum effect, it can never be fully eliminated, however by increasing the amount of power circulating in the interferometer, it is possible to reduce the photon counting uncertainties, hence limiting the impact of the shot noise. A power recycling mechanism is typically deployed to increase the circulating power in the instrument, at the expense of an increased noise from radiation pressure.

- *Radiation pressure noise*, emerging from the laser radiation pressure on the mirror. As for the seismic noise, the radiation pressure noise acts on the low frequencies, however it cannot be mitigate by a suitable suspension system. In principle the pressure noise can be straightforwardly reduced with a low circulating power, but this would result in a high amount of shot noise at high frequencies. Theoretically, the trade off between radiation pressure noise and shot noise is a direct consequence of the Heisenberg's uncertainty principle and the optimal balance the two sources can be achieved by light squeezing [205, 206] or even by a frequency dependent light squeezing [207]. The future Einstein telescope will eventually go beyond this limitation by constructing two interferometers: one with a low circulating power optimized to reduce the pressure noise in the low frequencies and another one with high circulating power, designed to reduce the shot noise in the high frequencies.
- *Other sources of noise* are also present in the interferometer, such as noise arising from the vacuum, from imperfections in the mirrors, from power fluctuations or from unideal laser or electronic components [85].

The properties of the detector's output are studied and validated as soon as new data are recorded. If some of the detector's output does not meet the quality requirements, a data quality veto can be issued [208–210]. A veto is a label, applied to a certain segment of the data, warning the user of a potential non optimal behaviour of the detector. Depending on the type of analysis, a scientist can choose to remove certain segments, based on the assessed data quality.

All the sources of noise described above produce a *stationary* noise, whose average properties do not depend on time. Besides stationary noise, sometimes loud transient burst of noise (*glitches*) [211–214] are possibly recorded. Glitches can happen in each detector with a rate as high as 1 per minute and they are empirically grouped into different categories based on their morphology [212]. While the origin of some glitch classes is known [215], in most cases the physical mechanism producing a glitch is unknown [216]. Glitches are an unavoidable part of the detector's output and they can negatively affect the extraction of physical information from the interferometer data. Indeed, they might mimic a short duration BBH signal, hence increasing the false positive rate of a search for BBH signals (see Sec. 3.4.2), or they may overlap with an astrophysical signal, hence reducing the amount of information available, as was the case for the BNS event GW170817 [217, 218]. Due to their ubiquity and their negative impact on the data analysis, substantial effort has focused on glitch subtraction, characterization and modeling [219–231].

2.2 Noise Characterization

Noise is the predominant output of any present ground-based interferometer. Indeed with the current sensitivity, BBH signals are recorded only every few days and, even if they are present in detector's output, their amplitude is typically one or two orders of magnitude smaller than the noise. For this reason, characterizing the statistical properties of the noise is essential. This knowledge, combined with certain assumptions about the noise, plays a crucial role in every data analysis application, including the searches for precessing black hole binaries discussed in this work.

2.2.1 Essentials of Timeseries Analysis

Before diving into the details, we introduce some notation and define some useful operations. The detector's output consist of a timeseries $s(t)$, which provides a noisy estimation of the strain Eq. (2.6) as a function of time. Practically, the strain timeseries is recorded only at some discrete times t_n , separated by a constant spacing Δt . We call sampling frequency, or sampling rate, f_s the quantity $1/\Delta t$. The detector's output is sampled with $f_s = 16\,384$ Hz but for a typical data analysis it is sufficient to set $f_s = 4096$ Hz or even $f_s = 2048$ Hz². When a timeseries $s(t)$ of length T is sampled at a given sample rate, it can be represented by a real vector \mathbf{s} such that:

$$s_n = s(t_n) \quad \text{with } n \in \{0, \dots, N-1\} \quad (2.13)$$

where, of course, $N = \lfloor T/\Delta t \rfloor = \lfloor Tf_s \rfloor$ and $\lfloor \cdot \rfloor$ denotes the floor of a real number.

Sometimes it is convenient to represent a timeseries in the frequency domain. This amount to applying the Fourier transform \mathcal{F} to $s(t)$ to obtain a *complex* frequency series $\tilde{s}(f) = \mathcal{F}s(t)$

$$\tilde{s}(f) = \int_{-\infty}^{+\infty} dt s(t) e^{-i2\pi f t}. \quad (2.14)$$

Clearly, the transformation is invertible:

$$s(t) = \int_{-\infty}^{+\infty} df \tilde{s}(f) e^{i2\pi f t}. \quad (2.15)$$

Note that if $s(t)$ is real, then $\tilde{s}(f) = \tilde{s}^*(-f)$. This means that is there is a symmetry between the real and the negative frequencies of the frequency domain timeseries.

²Indeed a typical BBH signal does not span frequencies higher than a thousands of Hertz, which is below the maximum measurable frequency at a given sample rate (also known as the Nyquist frequency). Moreover, as discussed in Sec. 2.1.2, the detector sensitivity is severely limited by the shot noise for frequencies above a few thousands Hertz.

For a discretely sampled timeseries, the Fourier transform reduces to the discrete Fourier transform (DFT), which produces a discrete frequency series:

$$\tilde{s}_k = \sum_{n=0}^{N-1} s_n e^{-i2\pi \frac{n}{N} k} \quad (2.16)$$

with inverse

$$s_n = \frac{1}{N} \sum_{k=0}^{N-1} \tilde{s}_k e^{i2\pi \frac{k}{N} n}. \quad (2.17)$$

The discrete frequency series \tilde{s} consist of N points and it is evaluated on a frequency grid with a constant spacing $\Delta f = \frac{1}{T}$:

$$f_k = \{-f_{\text{Ny}}, -f_{\text{Ny}} + \Delta f, \dots, 0, \dots, f_{\text{Ny}} - \Delta f, f_{\text{Ny}}\} \quad (2.18)$$

where $f_{\text{Ny}} = \frac{1}{2} f_s = \frac{1}{2\Delta t}$ is known as the Nyquist frequency and it amounts to the maximum measurable frequency given the time resolution available.

For most operations, included spectral estimation and whitening described below, it is useful to apply a window to the data, i.e. to multiply the timeseries $s(t)$ by a window function $W(t)$. The window function is usually taken to have zero values at the edges of the time grid. Windowing a timeseries avoids the edge-effects due to the limited sample size, arising from computing the discrete Fourier transform.

Finally, we define the cross-correlation \star between two timeseries $a(t), b(t)$ as:

$$(a \star b)(t) = \int_{-\infty}^{\infty} d\tau a(\tau - t)b(\tau). \quad (2.19)$$

The cross-correlation has a remarkable property: the Fourier transform of the cross-correlation of two timeseries is the product of the Fourier transform of the two timeseries - in symbols $\mathcal{F}\{a \star b\} = \mathcal{F}\{a\}\mathcal{F}\{b\}$. This is the analogous of the convolution theorem for cross-correlation. The operation of cross-correlation lies at the core of the matched-filtering algorithm and it is extensively used in the search for BBH signals, as detailed in Sec. 3.1.

2.2.2 A statistical Model for the Noise

Since, for every given time, the noise part n_t of the detector's output Eq. (2.12) takes a random value, we can conveniently model the noise timeseries n_t as a *stochastic process*. A stochastic process is a collection of random variables indexed by a variable t , which may be correlated with each other. For a stochastic process, we can define the mean $\mu(t)$ and the two-point correlation function $K(t_1, t_2)$

$$\mu(t) = E[n_t] \quad (2.20)$$

$$K(t_1, t_2) = E[n_{t_1} n_{t_2}] \quad (2.21)$$

where the expected value E is taken over many noise realizations in a procedure called *ensemble average*.

In what follows, we *assume* that the noise process n_t in consideration is fully characterized by its mean and its two-point correlation function. This means that other quantities derived from the process, such as the three-point or higher-order correlations, are functions of these two quantities. It turns out [232,233] that our knowledge of $\mu(t)$ and $K(t_1, t_2)$ is sufficient to write down the probability distribution p for a realization of the noise (n_0, \dots, n_{N-1})

$$p(n_0, \dots, n_{N-1}) \propto \exp \left\{ -\frac{1}{2} \sum_{ij} (n_i - \mu_i) K_{ij}^{-1} (n_j - \mu_j) \right\} \quad (2.22)$$

where $K_{ij} = K(t_i, t_j)$, $\mu_i = \mu(t_i)$ and K^{-1} denotes matrix inversion. The result is obtained by applying the Maximum Entropy principle, which states that such distribution must maximise the entropy, compatibly with the available constraints given by the knowledge of $\mu(t)$ and $K(t_1, t_2)$. The solution to the maximization problem is then a simple multivariate Gaussian³. For this reason, a process satisfying Eq. (2.22) is called *Gaussian noise*. It is important to recognize that the probability distribution Eq. (2.22) is derived solely from the knowledge of $\mu(t)$ and $K(t_1, t_2)$. Had we included additional information in our noise characterization (e.g. an expression for the three-point correlation), we would have obtained a different probability distribution for the observed timeseries.

A crucial observation is that the mean of the detector's noise is zero, i.e. $\mu(t) = 0$. Moreover, the two-point correlation function $K(t_1, t_2)$ is found to depend only on $\tau = t_1 - t_2$ and hence is invariant under a time shift $t_1, t_2 \rightarrow t_1 + C, t_2 + C$. A process satisfying this property is said *wide-sense stationary*. The empirical observation above is only valid for when averaging over a "short" period of time $O(\text{hours})$. However, as our analysis will focus on signals with a duration $O(10 \text{ seconds})$, we can safely assume that the interferometer noise is stationary over the time of analysis. The properties of stationary noise, i.e. $K(t_1, t_2)$, will still change over the timescale of hours and this will be taken into account. See [234–238] for some examples on how to drop this assumption. For a wide-sense stationary process, we can define the *autocorrelation function* $C(\tau)$:

$$C(\tau) = K(\tau, 0) = E[n_0 n_\tau]. \quad (2.23)$$

Thanks to our Gaussian noise assumption, the knowledge of the autocorrelation fully specifies the stochastic properties of the noise, computed by invoking the Maximum Entropy principle.

³The fact that the Gaussian distribution maximizes entropy given the variance and the mean is indeed a standard textbook results in statistics and it justifies its ubiquitous application throughout physics.

A deeper understanding on the properties of stationary noise arises from considering the correlation $E[\tilde{n}_f^* \tilde{n}_{f'}]$ of different frequencies, averaged over time:

$$E[\tilde{n}_f^* \tilde{n}_{f'}] = E \left[\int_{-\infty}^{+\infty} dt dt' n(t) n(t') e^{i2\pi f t} e^{-i2\pi f' t'} \right] \quad (2.24)$$

$$= \int_{-\infty}^{+\infty} dt d\tau C(\tau) e^{-i2\pi t(f' - f)} e^{-i2\pi f' \tau} \quad (2.25)$$

$$= \delta(f - f') \int_{-\infty}^{+\infty} d\tau C(\tau) e^{-i2\pi f' \tau} \quad (2.26)$$

$$= \delta(f - f') S(f) \quad (2.27)$$

where in the computation we used $C(t' - t) = E[n_0 n_{t' - t}]$, we performed the change of variables $\tau = t' - t$ and in the last line we defined:

$$S(f) = \int_{-\infty}^{+\infty} d\tau C(\tau) e^{-i2\pi \tau f}. \quad (2.28)$$

Note that as $C(\tau)$ is a real function, $S(f) = S(-f)$. Eq. (2.27) states that, for a wide-sense stationary process, the components in the frequency domain are uncorrelated with each other. Moreover, it provides a simple characterization of the variance $E[|\tilde{n}_f|^2]$ of the frequency domain noise, which simply amounts to the Fourier transform of the autocorrelation function.

Further insight is given by the Wiener-Khinchin theorem, which states that $S(f)$ is equivalent to the *power spectral density* (PSD), defined as the absolute value of the Fourier transform $\tilde{n}_T(f)$ of the noise timeseries $n_T(t)$ with length T in the limit of large values of T :

$$S(f) = \lim_{T \rightarrow \infty} \frac{1}{T} |\tilde{n}_T(f)|^2. \quad (2.29)$$

This expression gives a useful characterization of $S(f)$ as well as a way to estimate its value. The PSD has units of time and it is commonly measured in 1/Hz. It is sometimes useful to work with square root of the power spectral density, which goes under the name of amplitude spectral density (ASD). In the inset of Fig. 2.1, we show the typical ASD for the two LIGO detectors, as measured around the first detection of a gravitational wave [188].

As we did for the time domain case, we can write the maximum entropy probability density of obtaining a noise realization of frequency domain noise $(\tilde{n}_0, \dots, \tilde{n}_{N-1})$. By rearranging Eq. (2.22) and assuming zero mean, we obtain:

$$p(\tilde{n}_0, \dots, \tilde{n}_{N-1}) \propto \exp \left\{ -\frac{1}{2} \sum_i \Delta f \frac{|\tilde{n}_i|^2}{S(f_i)} \right\} \quad (2.30)$$

where the frequencies f_i belongs to the usual grid Eq. (2.18) and of course, $\Delta f = 1/T$. By taking the continuous limit, we obtain:

$$p[\tilde{n}(f)] \propto \exp \left\{ -\frac{1}{2} \int_{-\infty}^{+\infty} df \frac{|\tilde{n}(f)|^2}{S(f)} \right\}. \quad (2.31)$$

This simple expression can be used both for sampling noise realizations and to compute the *likelihood* that a given frequency series $\tilde{n}(f)$ is generated by the noise process described by the power spectral density $S(f)$.

Motivated by Eq. (2.31), we introduce the Wiener complex scalar product $\langle \cdot | \cdot \rangle$ between two timeseries $a(t), b(t)$ as follows:

$$\langle a | b \rangle = \int_{-\infty}^{+\infty} df \frac{\tilde{a}^*(f) \tilde{b}(f)}{S(f)}. \quad (2.32)$$

We denote with $\langle \cdot | \cdot \rangle$ and $[\cdot | \cdot]$, the real and imaginary part respectively of $\langle \cdot | \cdot \rangle$, so that

$$\langle a | b \rangle = (a | b) + i[a | b]. \quad (2.33)$$

Using this notation, the probability distribution for Gaussian noise takes a simple expression:

$$p[\tilde{n}(f)] \propto e^{-\frac{1}{2} \langle n | n \rangle}. \quad (2.34)$$

We use the $\hat{\cdot}$ symbol to denote a normalized timeseries

$$\hat{a} = \frac{1}{\sqrt{(a | a)}} a(t) \quad (2.35)$$

such that $(\hat{a} | \hat{a}) = 1$.

In a practical application, the strain is only calibrated for a limited range of frequencies and for this reason, it is often convenient to consider only a certain range of positive frequencies $[f_{\min}, f_{\max}]$, where f_{\min} and f_{\max} depend on the application. We can also introduce the *one sided* power spectral density $S_n(f)$, defined only for positive frequencies $f > 0$ as:

$$S_n(f) = 2S(f). \quad (2.36)$$

By setting a frequency range and using the one sided PSD, we can re-define the Wiener product as

$$\langle a | b \rangle = 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{a}^*(f) \tilde{b}(f)}{S_n(f)} \quad (2.37)$$

This definition allows the scalar product to be computed only within a specific frequency window of interest. Typically, we set $f_{\min} = 10/15$ Hz and $f_{\max} = 1024/2048$ Hz, thereby excluding frequency ranges where the detector's sensitivity is poorer (see the inset of Fig. 2.1). In this work, we will consistently apply this definition for the Wigner product, and when relevant, we will explicitly specify the frequency range considered. Clearly, Eq. (2.34) is still valid within the domain of interest.

2.2.3 Spectral Estimation

From the discussion above, it is clear that computing the power spectral density is of primary importance for the characterization of noise. While several spectral estimation methods have been developed, gravitational-wave data analysis mostly relies on the *Welch's* method [239]. The method relies on Eq. (2.29) and computes the discrete Fourier transform of a noise realization to estimate the PSD. To reduce the bias of the estimation, the time domain signal is usually sliced into several smaller segments: each segment is windowed and after computing the DFT, the results are averaged together. While the averaging procedure reduces the estimation bias for an individual point, it downgrades the method resolution, especially at the low frequencies, due to less frequency points considered.

Even though the method is widely used, mostly due to its simplicity, the trade off between resolution and bias requires a lot of tuning and the PSD estimation can suffer from a poor choice of the length of segments. Also the windowing procedure involves some arbitrary choices which may impact negatively the PSD estimation.

Maximum entropy spectral estimation A competitive alternative to the Welch's method comes from applying the maximum entropy principle to the problem, as first done by Burg [240]. The method was recently brought to the attention of the gravitational-wave data analysis community in [9], with some interesting application to parameter estimation.

The starting point of the Burg's method is the computation of the N empirical autocorrelation coefficients \hat{C}_k , which provides an estimator of the autocorrelation $C(\tau = t_k)$ based on a *single* noise realization x_t . For an infinitely long timeseries, the empirical autocorrelation coefficients are known for any discrete time shift t_k , hence the problem of computing the spectrum reduces to just computing the discrete Fourier transform of \hat{C}_k , by evaluating Eq. (2.28). However, due to the finite size of the sample, our knowledge of such coefficients is limited. Moreover, the variance of their estimation will increase as the time shift increases, further limiting our knowledge of the autocorrelation. As a way out, we can invoke the maximum entropy principle to find the PSD as the solution of a variational problem which takes into account our knowledge and at the same time *does not* make any assumption on the unknown coefficients. The variational problem can be stated as follow. We need to find a function $S(f)$ which maximizes the entropy gain ΔH

$$\Delta H = \int_{-f_{Ny}}^{f_{Ny}} df \log S(f) \quad (2.38)$$

and whose Fourier transform agrees with the empirical autocorrelation coefficients \hat{C}_k

$$\int_{-f_{Ny}}^{f_{Ny}} df S(f) e^{i2\pi fk\Delta t} = \hat{C}_k. \quad (2.39)$$

Note that, unlike with the Welch method, the maximization problem does not make any assumption about the unknown coefficients C_k with $k > N$.

The maximization problem admits a closed form solution:

$$S(f) = \frac{P\Delta t}{|1 - \sum_{k=1}^p a_k e^{i2\pi fk\Delta t}|^2} \quad (2.40)$$

where the p coefficients a_k as well as the scale factor P are obtained by solving a recursive formula. The number p of coefficients is the only tunable parameter for the PSD estimation and several strategies are available for an effective choice.

A comparison between the Burg's algorithm and the Welch's algorithm reveals that the Burg's solution yields an estimation with a smaller variability over different noise realizations and, especially for short timeseries, a better accuracy for the PSD at low frequencies [9]. A PSD estimation stable for very short timeseries can open up interesting opportunities in the parameter estimation of a BBH signal, possibly avoiding any bias introduced by using the Welch's method.

2.2.4 Whitening

We saw that in Gaussian noise samples at different times are correlated with each other and a measure of this correlation is given by the autocorrelation function $C(\tau)$ or equivalently, by its Fourier transform, the power spectral density $S(f)$. For many applications, including searching for signals, it is useful to introduce a transformation, called *whitening*, which removes any correlation between different times. We therefore obtain an uncorrelated noise process $n^W(t)$, called *white noise*, each sample is independent from the others, hence the autocorrelation function is simply $C(\tau) = \delta(\tau)$ and, according to Eq. (2.22), the noise at every time is sampled from a standard normal distribution (i.e. zero mean and unit variance):

$$n_t^W \sim \mathcal{N}(0, 1). \quad (2.41)$$

The whitening transformation acts on a Gaussian noise process $n(t)$ to produce a white noise process $n^W(t)$ as follows

$$n^W(t) = \int_{-\infty}^{+\infty} df \frac{\tilde{n}(f)}{\sqrt{S(f)}} e^{i2\pi ft} \quad (2.42)$$

A simple computation of the two-point correlation function shows that indeed Eq. (2.42) produces a white noise process:

$$\mathbb{E} [n_{t_1}^W n_{t_2}^W] = \int_{-\infty}^{+\infty} df df' \frac{\mathbb{E} [\tilde{n}^*(f) \tilde{n}(f')]}{S(f)} e^{i2\pi f t_1} e^{-i2\pi f' t_2} \quad (2.43)$$

$$= \int_{-\infty}^{+\infty} df e^{i2\pi f(t_1 - t_2)} = \delta(t_1 - t_2). \quad (2.44)$$

where we used Eq. (2.27) to compute the expected value over different noise realizations. Finally, we can see that the PSD of white noise process is constant

$$S^W(f) = \int_{-\infty}^{+\infty} d\tau \delta(\tau) e^{-i2\pi f\tau} = 1. \quad (2.45)$$

Clearly, any white noise process can be rescaled by a constant factor to set a given variance. This rescaling also affects the PSD.

As introduced above, the whitening operation removes the correlation between samples at different times, which is expected from a Gaussian noise model for the timeseries. This step is crucial because it simplifies the detection of any residual correlation, which may arise either from an astrophysical signal or a non-Gaussian noise transient. Further insight can be obtained by looking at the frequency domain representation of the noise timeseries. As expressed in Eq. (2.45), whitening rescales each frequency so that the variance of the power in each frequency bin becomes constant. In this way, the whitening process normalizes the “noise floor” across different frequencies, making it easier to detect excess power in any particular frequency bin. An excess power in a frequency bin manifests in the time domain as a correlation between samples at different times, highlighting the complementary nature of the time and frequency domain perspectives. This relationship underscores how insights from both domains can be combined to improve signal detection and noise characterization.

2.3 Overview of the Data Analysis for Compact Binary Coalescences

An array of interferometers produces a large amount of raw data, which need to be analysed with complex techniques in order to extract information about compact binaries and answer the most pressing questions of gravitational-wave astronomy. We discuss below a few of such techniques, highlighting some of the questions that the gravitational-wave community is trying to address. Far from being exhaustive, our discussion is helpful to frame the current work about searching for precessing compact binary signals into the large and active field of gravitational-wave astronomy.

Note that our discussion only focuses on the physics of compact binary coalescence (CBC), hence neglecting a large portion of the scientific goals of the LIGO, Virgo and KAGRA (LVK) collaboration, such as studies in continuous waves [241–245], supernovae [246–250] and stochastic gravitational-wave background [251–255].

Waveforms A preliminary step for *every* data analysis application on CBC study is the availability of fast and accurate prediction of the gravitational signal emitted by a BBH, a binary neutron star (BNS) or a neutron star-black hole system (NSBH). As already discussed in some details in Sec. 1.3.2, this is done by “combining” the perturbative dynamics of the binary with accurate but expensive predictions from numerical relativity. This idea gave birth to two families of model, EOB models [52, 153, 155, 156, 256] and phenomenological model [48, 157, 257–259], both readily available for the public and widely used. Such models may or may not incorporate effects of spins, spins precession, higher order multipoles and (more recently) eccentricity.

As most of the data analysis relies on a fast availability of such waveforms, a longstanding effort has been devoted to develop *surrogate* models [260–267], with the goal of speeding up the waveform generation and enable several data analysis applications. Surrogate models are usually large parametric models, trained to reproduce waveforms from an approximant of choice or even by numerical relativity waveforms. Recent works have developed machine learning based surrogate models [5–7, 268–276], even though none of such model has yet reached a production level.

Searches Once the interferometer data are available, the first task is to identify a set of candidate CBC signals and assign them a measure of the statistical significance of a potential detection. This is the primary goal of a search pipeline, which can be broadly categorized into unmodeled and modeled. Unmodeled searches, such as those of the Coherent WaveBurst (cWB) pipeline [277–281], search for transient signals with minimal assumptions by looking for coherent excess power between two detectors. Unmodeled searches tend to be less sensitive to CBC signals than modeled searches but can potentially identify signals from a broader range of sources, even the ones that the theory fails to predict. On the other hand, modeled searches rely on matched-filtering [63] to correlate the detector’s output with a large number of CBC waveform templates. While they generally offer greater sensitivity compared to unmodeled searches, this comes at the cost of reduced sensitivity to signals outside their target regions or to other non-CBC events. We will discuss matched-filtering searches in the next Chapter. Both types of searches have been used to build most recent LVK GW transient catalogs [15–18] and they were crucial to detect all the confirmed GW events detected so far.

Parameter estimation Once a waveform is detected, a thorough analysis is needed to provide a measurement of the properties (masses, spins, sky location etc...) of the source binary. This is done by means of a *parameter estimation* procedure [282–284] which operates in a Bayesian framework. Given a set of parameters θ of a parametric model $p(\mathcal{D}|\theta)$ which describes the observed data \mathcal{D} , the analysis defines a posterior probability distribution $p(\theta|\mathcal{D})$, which represents our best knowledge of the source.

Sampling from the posterior $p(\theta|\mathcal{D})$ is a challenging task, due to the high dimensionality of the problem and to high computational cost related to the generation of millions of waveforms. Traditionally the parameter estimation utilizes standard sampling techniques, such as Markov chain Monte Carlo [285] and Nested Sampling [286, 287]. More recently Machine Learning based sampling approaches have been investigated [288–292].

Rates and population The combined detections of many CBC sources provide invaluable information about the population of astrophysical binaries, its properties and the rate of mergers per unit of time and volume. Several methods have been developed to estimate the rate of a population of binaries based on search results [293, 294]. Moreover, ever since tens of detections are routinely performed, more sophisticated Bayesian methods have emerged not only to estimate the signal rates but also to constrain the statistical distributions of astrophysical parameters such as masses, redshifts and spins [19–21]. To achieve this, both parametric [295–297] and non parametric [298–302] probabilistic models have been developed and are routinely used. These models rely on the results of parameter estimation for each detected event to make the most accurate inference about the population’s properties.

Tests of general relativity A detected CBC signal provides a unique regime where to test the predictions of theory of GR in a strong field regime [303]. By comparing the many detected signals with the GR predictions, it is possible to probe the theory in many scenarios, allowing to rule out alternative theories and to pose experimental constraints on the amount of GR violation. Tests of general relativity routinely performed include constraints on the dynamics of the inspiral or on the ringdown phase [304, 305], a measurement of the speed of propagation of a GW and a measurement of the number of polarizations of the spacetime perturbations [30–34]. So far, no deviation from the theory of general relativity has been observed.

Extreme matter and multi-messenger astronomy A binary neutron star merger [306, 307] is a unique phenomenon where matter is squeezed at very high densities $\rho \simeq 10^{14} \text{ g/cm}^3$, in such extreme conditions that the macroscopic and microscopic behaviour of matter are still poorly understood. A GW from a binary neutron star sys-

tem can give invaluable information on the deformability of a neutron star [308], allowing to pose important constraints on the equation of state the matter at high densities. Moreover, a binary neutron star merger will launch an energetic gamma ray burst [309, 310], with consequent ejection of highly radioactive matter in what is known as a kilonova [311]. Both the gamma ray burst and the kilonova can be observed in the electromagnetic band with existing telescopes. The multi-messenger study of binary neutron stars has started with the observation GW170817 [23, 217, 312], with plenty of spectacular results [22, 24, 25, 31, 313–315], and has the potential of providing more interesting observations, as the interferometers collect more data.

Cosmology The gravitational wave from a compact binary allows for a measurement of the distance of the binary. This information, coupled with a measure of redshift, can be used to constraint the expansion history of the universe, via the measurement of the Hubble constant H_0 [26, 28, 29, 316]. A few methods have been developed for redshift estimation [317, 318], either jointly fitting the cosmological models and the BBH redshift distribution or using statistical information from galaxy catalogs. In the case of a multi-messenger event, such as GW170817, the redshift can be easily obtained by the electromagnetic observation of kilonova surrounding the merger [319, 320]. While the uncertainties for the inferred value for the Hubble constant are still large, more observation might eventually solve the “Hubble tension” [321] between the measurements of the H_0 from cosmic microwave background [322] and from standard candles [323, 324].

CHAPTER 3

Searching for Signals with Matched Filtering

In this section we describe some general features of the matched-filtering pipelines that are routinely employed to search for binary black holes systems. Their primary objective is to produce a list of potential candidates, ranked from the most to the least likely to be of astrophysical origin, and to estimate the statistical significance of each candidate. We will first introduce the basics of matched filtering and show how it can be used to define a suitable search statistic useful to identify potential signals. The search statistic must be evaluated for many different binary systems of interest, forming large template banks, whose generation will be the topic of one subsection. Finally, we will provide an overview of a realistic matched-filtering pipeline, diving into some details of the several steps required to build a candidate list and to estimate their statistical significance. We will mostly specialize our discussion to the GstLAL pipeline, even though we will highlight differences and features with other widely used pipelines.

3.1 Basics of Matched Filtering

A major challenge for the detection of gravitational waves comes from the fact that the magnitude of an *oscillating* signal $h(t)$ is one or two orders of magnitude smaller than the noise $n(t)$. Luckily, in this scenario it is still possible to detect a signal thanks to the technique of matched filtering. Heuristically, matched filtering computes the correlation $\frac{1}{T} \int_0^T dt h(t)s(t)$ between a signal $h(t)$, which starts at time $t = 0$ and has

a duration T , and the noisy data $s(t) = n(t) + h(t)$:

$$\frac{1}{T} \int_0^T dt h^2(t) + \frac{1}{T} \int_0^T dt h(t)n(t). \quad (3.1)$$

The advantages of this expression are evident: while the first term is constant and amounts to the squared amplitude of the signal¹, the second term decreases as² $1/\sqrt{T}$. As a result, for sufficiently long (or loud) signals, this method can potentially reveal their presence even when buried deep within the noise floor.

In the rest of the section we will make this rough argument more precise, employing two different approaches to obtain the expression for the optimal filter to be used to search for a signal $h(t)$. A matched-filtering search will then compute the correlation between the data $s(t)$ and the optimal filter as a function of time (i.e. as a function of a time shift of the template). Consequently, it will claim a detection only if the correlation value, also called *search statistic*³ or Signal-to-Noise ratio, exceeds a certain threshold at a given time.

3.1.1 Matched Filtering as Optimal Filter

To formalize our problem [14, 85, 325], we define the correlation ρ between the detector's output $s(t)$ and a filter $K(t)$

$$\rho = \int_{-\infty}^{\infty} dt s(t)K(t). \quad (3.2)$$

This is real random variable, whose value depends on a particular noise realization and on the signal $h(t)$ present in the detector's data. We want to find the *optimal filter* $K(t)$ so as to maximise the ratio S/N between the expected value S of ρ when a signal is present and the square root N of the variance of ρ in the presence of noise only. After the optimal filter is chosen, the value of ρ represents the search statistic used to claim a detection.

In the presence of a signal, the expected value S of ρ is

$$S = E \left[\int_{-\infty}^{\infty} dt s(t)K(t) \right] = \int_{-\infty}^{\infty} df \tilde{K}^*(f)\tilde{h}(f), \quad (3.3)$$

¹For a perfect sine wave, the integral $\int_0^T dt h^2(t)$ is equal to Th_0^2 , where h_0 is the amplitude of the wave and therefore $\frac{1}{T} \int_0^T dt h^2(t) = h_0^2$. While this argument relies on a sinusoidal wave, it can be easily extended to a chirp-like wave, such as a gravitational-wave signal. Indeed, we could show that for any oscillating signal, the integral $\int_0^T dt h^2(t)$ is always proportional to T .

²As h and n are uncorrelated, $\int_0^T dt h(t)n(t)$ is well approximated by the integral of a random walk process, whose magnitude grows as \sqrt{T} .

³As explained below, this term is borrowed from the field of frequentist hypothesis testing.

where we used the fact that $E[s(t)] = E[h(t)] = h(t)$ and we wrote the convolution in the frequency domain. On the other hand, the variance N^2 of ρ when a signal is absent is given by

$$N^2 = E[\rho^2] - E[\rho]^2 \quad (3.4)$$

$$= \int_{-\infty}^{\infty} df df' \tilde{K}^*(f) \tilde{K}(f') E[\tilde{n}^*(f) \tilde{n}(f')] \quad (3.5)$$

$$= \int_{-\infty}^{\infty} df |\tilde{K}(f)|^2 S(f) \quad (3.6)$$

where we used the fact that $E[\rho] = 0$ and we used Eq. (2.27) to evaluate $E[\tilde{n}(f) \tilde{n}(f')]$.

The ratio S/N is usually called Signal-to-Noise ratio (SNR) and can be conveniently written using the real part of the Wiener product Eq. (2.32):

$$\text{SNR} = \frac{S}{N} = \frac{(u|h)}{\sqrt{(u|u)}} = (\hat{u}|h) \quad (3.7)$$

with $\tilde{u}(f) = S(f) \tilde{K}(f)$ and $\hat{u} = \frac{u}{\sqrt{(u|u)}}$ as in Eq. (2.35). Trivially, the expression is maximized when $u \propto h$, hence the optimal filter to search for a signal $h(t)$ in the detector's data is

$$\tilde{K}(f) \propto \frac{\tilde{h}(f)}{S(f)} \quad (3.8)$$

where an overall constant can be freely chosen. By combining the last two equations, we obtain an expression for the SNR of a signal $h(t)$

$$\text{SNR}^2 = (h|h) = \Re \int_{-\infty}^{+\infty} df \frac{|\tilde{h}(f)|^2}{S(f)}. \quad (3.9)$$

With the optimal filter, we can then write down the final expression for search statistic ρ :

$$\rho = (s|\hat{h}) = \Re \int_{-\infty}^{\infty} df \frac{\tilde{s}^*(f) \tilde{h}(f)}{S(f)}, \quad (3.10)$$

where we set to one the amplitude of the optimal filter in Eq. (3.8). This expression is completely generic and does not make any assumption about the nature of the signal $h(t)$ we are searching for. In the next section, we will specialize the equation to the case of a compact binary coalescence.

3.1.2 Matched Filtering as Hypothesis Test

Another, perhaps more elegant, approach for determining the optimal search statistic comes from the theory of hypothesis testing [141]. In this context, we frame the detection problem as a test between two hypothesis:

1. *The “noise” hypothesis*: The observed strain $s(t)$ is only formed by noise
2. *The “signal” hypothesis*: The observed strain $s(t)$ is formed by a superposition of noise and a signal $h(t)$

We then need to perform a test to distinguish between the two hypothesis. In a frequentist framework, this involve the computation of a *test statistic*, whose value determines whether the “signal” hypothesis can be accepted. According to the Neyman–Pearson lemma [326], the likelihood ratio between the two hypothesis is the optimal score for the hypothesis test in the sense that it maximizes the true positive rate, i.e. the probability of accepting a signal hypothesis whenever a signal is actually present in the data.

Mathematically, we can assign a likelihood to the “noise” hypothesis by evaluating the probability that the observed strain arises from Gaussian noise, as given in Eq. (2.34). The likelihood for the “signal” hypothesis can be obtained by considering that if the signal $h(t)$ is present in the data $s(t)$, then $s(t) - h(t)$ must be a timeseries formed of noise only and should also follow Eq. (2.34). The log likelihood ratio Λ is then

$$\Lambda = \log \frac{p(s(t)|\text{signal})}{p(s(t)|\text{noise})} = \log \frac{e^{-\frac{1}{2}(s-h|s-h)}}{e^{-\frac{1}{2}(s|s)}} = (s|h) - \frac{1}{2}(h|h), \quad (3.11)$$

where we used the standard notation $p(A|B)$ to denote the conditional probability of A given B . The conditional probability $p(A|B) = \frac{p(A \cup B)}{p(B)}$ amounts to the probability of the event A given the knowledge that the event B has already occurred.

We note that the log likelihood ratio Λ depends on the data only through the quantity $(s|h)$ and it is a monotonically increasing function of this inner product. This means that we are free to choose the amplitude of $h(t)$ without affecting the ability of the test to distinguish between signal and noise⁴. By convention, we choose such amplitude so as to maximise the value of Λ . The maximization over a constant scale factor a can be written as

$$\max_a \Lambda = \max_a \left\{ a(s|h) - \frac{a^2}{2}(h|h) \right\} \quad (3.12)$$

and it yields the expression for the search statistic $2\Lambda = \rho^2 = \frac{(s|h)^2}{(h|h)} = (s|\hat{h})^2$, identical to Eq. (3.10).

3.2 Search Statistic

The search statistic ρ in Eq. (3.10) is useful to detect a *specific* signal $h(t)$ within the data, also known as the *template*. The analytical form of a template must be known

⁴This is true provided that we scale the detection threshold accordingly.

with precision ahead of the search and it usually depends on many parameters stored in a vector ϑ , which describes the property of the source and fully specify a template. We may then write $h(t; \vartheta)$. Clearly, in a realistic scenario, we are not interested in checking for the presence of a single signal, characterized by a particular value of ϑ , but we rather seek to detect *any* signal present in the data, within a reasonable range of template parameters. For this reason, our search problem can be cast into a large optimization problem, where the search statistic ρ is maximized over all the possible values of the template parameters ϑ :

$$\max_{\vartheta} \rho(\vartheta) = \max_{\vartheta} (s | \hat{h}(\vartheta)). \quad (3.13)$$

The search statistic ρ Eq. (3.10) is fully general and does not make any assumption about the nature of signal or even about the type of detector. Without further knowledge of the signal model, the maximization Eq. (3.13) must be carried out by a brute force approach, where $\rho(\vartheta)$ is evaluated over a discrete set of value of ϑ , or equivalently, different matched-filtering searches are carried out for many different templates. The discrete set of template searched for is called *template bank*, which will be the topic of the next section.

By including the effect of eccentricity, the template h for a BBH signal is described by 17 parameters, as shown in Eq. (2.6) and Eq. (1.62). The two masses m_1, m_2 and spins s_1, s_2 of the BHs together with the eccentricity e of the orbit and the mean periastron anomaly a form a set of 10 parameters called *intrinsic*, since they only depend on the source. Other 7 parameters are needed to describe the position and orientation with respect to the observer and for this reason are called *extrinsic*: they are the time of arrival in the detector frame t_0 , the inclination angle ι , the reference phase ϕ , the polarization angle Ψ and the sky location, parameterized by the right ascension α , the declination δ , together with the luminosity distance D_L . Luckily, we can maximise the search statistic over some of these nuisance parameters, thus reducing the number of parameters that needs to be discretely covered by the template bank, i.e. the *dimensionality* of the bank. The details of the analytical maximization depend on the type of binary considered.

As a realistic search does not have access to all the frequencies components of the strain, every matched-filtering analysis is restricted only to a finite range of frequencies $[f_{\min}, f_{\max}]$. From a practical point of view, this amounts to computing the scalar product using the one-sided PSD $S_n(f)$ and to constrain any integration over frequencies to the range of interest. Thus, in what follows, we will consider the scalar product Eq. (2.37), instead of the scalar product Eq. (2.32) used above.

3.2.1 Statistic for Circular Aligned-spin Binaries

A standard assumption is that we only search for circular aligned-spin binaries, where only the quadrupole dominant $(\ell, |m|) = (2, 2)$ modes are considered. From an astrophysical perspective, restricting the search to circular, aligned-spin systems still encompasses a large fraction of the binaries in the universe. Indeed, most binaries are expected to have lost their orbital eccentricity through gravitational radiation by the time they merge, and the dominant formation channel—isolated binary evolution—is believed to produce systems with nearly aligned spins. Furthermore, this choice drastically simplifies the morphology of the waveform and by limiting the template bank size, it reduces the cost of the search by orders of magnitude, as we will see below with more details. Therefore, most of the traditional searches introduced so far are restricted to aligned-spin binaries with only the dominant mode considered.

Under this simplification, Eq. (1.62) becomes:

$$h_+ + i h_\times = \frac{1}{D_L} \left\{ h_{22} Y_{22}(\iota, \phi) + h_{2-2} Y_{2-2}(\iota, \phi) \right\} \quad (3.14)$$

$$= \frac{1}{D_L} \sqrt{\frac{5}{64\pi}} (1 + \cos \iota)^2 A e^{i\varphi} e^{2i\phi} + \frac{1}{D_L} \sqrt{\frac{5}{64\pi}} (1 - \cos \iota)^2 A e^{-i\varphi} e^{-2i\phi} \quad (3.15)$$

$$= \frac{1}{D_L} \sqrt{\frac{5}{64\pi}} \frac{1 + \cos^2 \iota}{2} A \cos(\varphi + 2\phi) + i \frac{1}{D_L} \sqrt{\frac{5}{64\pi}} \cos \iota A \sin(\varphi + 2\phi) \quad (3.16)$$

where we use the explicit expression for the spherical harmonics [103] and A and φ are the time-dependent amplitude and phase of the $(\ell, |m|) = (2, 2)$ modes

$$h_{22} = h_{2-2}^* = A e^{i\varphi}, \quad (3.17)$$

which depend only on the masses m_1, m_2 and on the z-components of the spins χ_{1z} and χ_{2z} ⁵. The expression is remarkably simple and extends the Newtonian formulas Eqs. (1.74-1.75) to account for a more accurate description of the orbital evolution.

The simplicity of Eq. (3.16) can be used to maximise the detection statistic Eq. (3.10) with respect to all the extrinsic parameters except the masses and spins, thus leaving only 4 parameters to maximise over by brute force. To show this, we note that the

⁵As we are focusing on circular binaries, we are not considering the eccentricity and the mean anomaly of the system.

strain h Eq. (2.6) can be written as [327]

$$h = F_+ h_+ + F_\times h_\times = \frac{1}{D_{\text{eff}}} A \cos(\varphi + 2\phi_0) \quad (3.18)$$

where the effective distance D_{eff} is given by

$$D_{\text{eff}} = \frac{D_L}{\sqrt{F_+^2 \left(\frac{1+\cos^2 \iota}{2}\right)^2 + F_\times^2 \cos^2 \iota}} \quad (3.19)$$

and the phase ϕ_0 is defined as

$$e^{2i\phi_0} = e^{2i\phi} \frac{D_{\text{eff}}}{D_L} \left[F_+ \frac{1 + \cos^2 \iota}{2} - i F_\times \cos \iota \right]. \quad (3.20)$$

This result is remarkable: the effect of all the extrinsic parameter on the observed signal can be absorbed by an overall amplitude scale and by a phase shift. This has a profound implication for matched-filtering, since we can easily maximise analytically the search statistic over an amplitude and a phase.

To move forward it is convenient to define the waveforms $h_p(t; m_1, m_2, s_{1z}, s_{2z})$, $h_c(t; m_1, m_2, s_{1z}, s_{2z})$ as

$$\begin{aligned} h_p(t) &= A(t) \cos(\varphi(t)) \\ h_c(t) &= A(t) \sin(\varphi(t)) \end{aligned} \quad (3.21)$$

with the notable properties that $h_p \propto h_+$, $h_c \propto h_\times$ and

$$\tilde{h}_p = i \tilde{h}_c. \quad (3.22)$$

Moreover, we note that h_p and h_c are orthogonal, i.e. $(h_c | h_p) = 0$. With this definition, the frequency domain strain can be written as:

$$\tilde{h} = \frac{1}{D_{\text{eff}}} \left[\tilde{h}_p \cos(\phi_0) + \tilde{h}_c \sin(\phi_0) \right] = \frac{1}{D_{\text{eff}}} e^{i\phi} \tilde{h}_p. \quad (3.23)$$

As the detection statistic does not depend on an overall amplitude scaling, we must only maximise over a constant phase ϕ_0 :

$$\max_{D_L, \alpha, \delta, \iota, \phi, \Psi} \rho^2(\vartheta) = \max_{\phi_0} \rho^2(\vartheta) = \max_{\phi_0} (s | \hat{h}_p e^{i\phi_0}). \quad (3.24)$$

The solution is trivial, as the effect of the phase is just to rotate the number $\langle s | \hat{h}_p \rangle$ on the complex plane. Thus the maximum value of the search statistic is attained when $\langle s | \hat{h}_p \rangle e^{i\phi_0}$ is a real number, i.e. when $\langle s | \hat{h}_p \rangle e^{i\phi_0} = |\langle s | \hat{h}_p \rangle|$, giving:

$$\max_{D_L, \alpha, \delta, \iota, \phi, \Psi} \rho^2(\vartheta) = |\langle s | \hat{h}_p \rangle|^2. \quad (3.25)$$

Thanks to the symmetry of the problem Eq. (3.22), the expression is equivalent to

$$\max_{D_L, \alpha, \delta, \iota, \phi, \Psi} \rho^2(\vartheta) = (s|\hat{h}_p)^2 + (s|\hat{h}_c)^2. \quad (3.26)$$

This result is straightforward to interpret: in the limit where only the dominant mode is considered, Eq. (3.26) allows to search for all the aligned-spin BBHs characterized by the same BH masses and spins, regardless of the binary orientation, position in the sky or choice of polarization. With this result, the template bank has to cover only a four dimensional space as opposed to a ten dimensional space that would result from using Eq. (3.10). This comes at the minimal price of filtering the data with the two templates \hat{h}_p and \hat{h}_c .

3.2.2 Statistic for General Binaries

The analytical maximization of the search statistic over the *extrinsic* parameters was made possible thanks to heavy assumptions about the nature of the signal. In the case where we don't make any of those assumptions, we can still analytically maximize over some nuisance parameters [328]. To do so, we start again from the strain h Eq. 2.6 and we rewrite it as:

$$h(t) = F_+ h_+(t) + F_\times h_\times(t) = \mathcal{A} \left[u \hat{h}_+(t) + \hat{h}_\times(t) \right] \quad (3.27)$$

where we defined the two variables

$$\mathcal{A} = F_\times \sqrt{(h_\times|h_\times)} \quad (3.28)$$

$$u = \frac{F_+}{F_\times} \sqrt{\frac{(h_+|h_+)}{(h_\times|h_\times)}}. \quad (3.29)$$

We are then able to maximise over \mathcal{A} and u , which capture the dependence of h on the sky position and on the polarization angle (4 parameters), leaving to a brute force approach the maximization over the remaining 10 parameters, i.e. the two BH masses, the two spin vectors \mathbf{s}_i , the inclination angle ι and the reference phase ϕ .

After a straightforward maximization over the amplitude \mathcal{A} , the maximization over u yields:

$$\max_{D_L, \alpha, \delta, \Psi} \rho^2(\vartheta) = \max_u \rho^2(\vartheta) = \frac{(s|\hat{h}_+)^2 + (s|\hat{h}_\times)^2 - 2\hat{h}_{+\times}(s|\hat{h}_+)(s|\hat{h}_\times)}{1 - \hat{h}_{+\times}^2} \quad (3.30)$$

where we define the important parameter $\hat{h}_{+\times}$ as

$$\hat{h}_{+\times} = (\hat{h}_+|\hat{h}_\times). \quad (3.31)$$

The quantity $\hat{h}_{+ \times}$ quantifies the lack of orthogonality between the two polarizations: while in the aligned-spin dominant mode case $\hat{h}_{+ \times} = 0$, this is not the case from precessing systems and/or systems showing imprints of higher order modes (HMs). We may informally refer to Eq. (3.30) as *symphony* search statistic, from the title of [328], which laid the foundation for our work. In [328] the term originally refers to the detection of higher-order modes, which combine multiple “voices” into a unified “symphony” produced by a BBH merger. Due to the similarities in their mathematical descriptions, we are going to use this term even if we consider precessing systems without imprints from higher order modes.

The expression can be simplified by introducing the *orthogonalized* template \hat{h}_\perp :

$$\hat{h}_\perp(t) = \frac{1}{\sqrt{1 - \hat{h}_{+ \times}^2}} \left(\hat{h}_\times(t) - \hat{h}_{+ \times} \hat{h}_+(t) \right) \quad (3.32)$$

which is orthogonal to \hat{h}_+ and properly normalized $(\hat{h}_\perp | \hat{h}_\perp) = 1$. With this definition, the search statistic takes a remarkably simple functional form:

$$\max_{D_L, \alpha, \delta, \Psi} \rho^2(\vartheta) = (s | \hat{h}_+)^2 + (s | \hat{h}_\perp)^2. \quad (3.33)$$

This is the same expression for the aligned-spin case Eq. (3.26), with the replacement $\hat{h}_p \rightarrow \hat{h}_+$ and $\hat{h}_c \rightarrow \hat{h}_\perp$. Note that while in both cases data are filtered with two orthogonal templates, the templates \hat{h}_+, \hat{h}_\perp depend on 10 parameters (two masses, six spins and the inclination angle and reference phase) while the aligned-spin templates \hat{h}_p, \hat{h}_c depend only on 4 parameters (two masses and two spins). For this reason, the complexity of the general case does not arise on a more complicated filtering expression but only on the higher computational cost introduced by a larger number of parameters to be spanned by a template bank. We can recover the aligned-spin case from Eq. (3.30), in the limit where $\hat{h}_{+ \times} \rightarrow 0$, where also $\hat{h}_+ \rightarrow \hat{h}_p, \hat{h}_\perp \rightarrow \hat{h}_c$.

We close by mentioning that, under the assumption that the BBH system is precessing but the signal is only composed by the dominant mode, we can further maximise over the reference phase ϕ [327, 329]. However, the resulting expression has a rather complicated functional form which i) increases the cost of filtering, ii) requires large changes to existing software infrastructures and iii) does not generalize to a search employing templates with HMs. Moreover, as we will see in Ch. 4, numerical studies have shown that the inclusion of the reference phase inside the template bank is not required to produce a bank of good quality. For all these reasons, we will not consider the search statistic in [327], but we will rather rely on Eq. (3.30) for the remainder of our work.

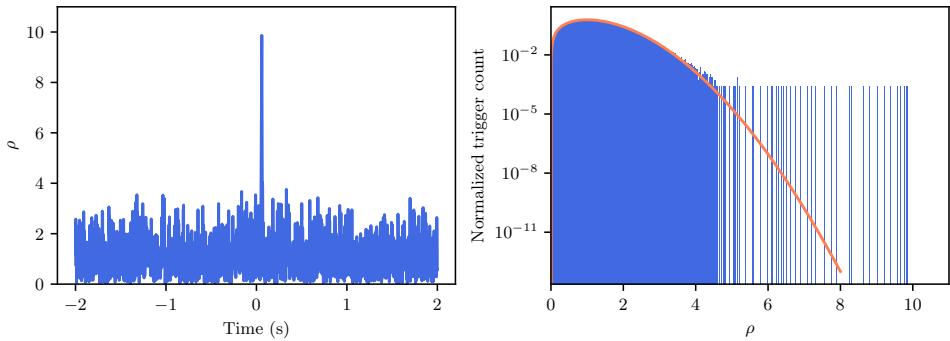


Figure 3.1: On the left plot we show the SNR timeseries $\rho(t)$ Eq. (3.35) for a non-precessing signal injected in Gaussian noise, filtered with a matching template. The signal is injected with $\text{SNR} = 10$ and refers to a system with $m_1 = 30 \text{ M}_\odot$ and $m_2 = 10 \text{ M}_\odot$, generated from a starting frequency $f_{\min} = 20 \text{ Hz}$. On the right plot, we show the normalized distribution of the measured values ρ of the search statistic, whose squares follow a chi-squared distribution with two degrees of freedom in presence of noise only. In orange, we overlay the chi-squared distribution for the values of ρ Eq. (3.41). Note that there are a few outliers to the SNR distributions, corresponding to the injected signal.

3.2.3 Maximization over Time

So far, we have described a method to search over different waveform parameters but we made the crucial assumption that the time of arrival t of the signal on the detector frame is known. Clearly, that is not the case and an appropriate search statistic must also maximise over the arrival time. This is done again by computing the search statistic for different times of arrivals to obtain a timeseries $\rho(t)$, which is the called the *SNR timeseries* with a slight abuse of notation. To keep our discussion general, we introduce two *real* templates⁶ h_R and h_I , which are not generally related to each other and are used to filter the data with the search statistic:

$$\rho^2 = (s|h_R)^2 + (s|h_I)^2 = \langle s|h_T \rangle^2 \quad (3.34)$$

where we introduced a complex template h_T , so that $h_T(t) = h_R(t) + i h_I(t)$. This case is very general and encompasses the search statistic maximized for both the situations described above. The aligned-spin search statistic Eq. (3.25) can be obtained by setting $h_R, h_I = \hat{h}_p, \hat{h}_c$ ⁷, while the general case Eq. (3.33) is retrieved with $h_R, h_I = \hat{h}_+, \hat{h}_\perp$.

⁶Note that here, to keep our discussion general, we are not making the assumption that the two templates are normalized.

⁷Thanks to the symmetry in Eq. (3.22), in frequency domain we have that $\tilde{h}_T = 2\tilde{h}_p$.

In the frequency domain, a constant time shift t of the template $h_T(f)$ is simply given by $h_T(f)e^{-i2\pi ft}$ and the SNR timeseries is readily computed from Eq. (3.34):

$$\rho(t) = \left| 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{s}^*(f) \tilde{h}_T(f)}{S_n(f)} e^{-i2\pi ft} \right| \quad (3.35)$$

where of course, the template \tilde{h}_T is normalized to unity. Motivated by this expression, we may also define a time dependent complex scalar product between the two frequencies $[f_{\min}, f_{\max}]$

$$\langle a|b \rangle(t) = 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{a}^*(f) \tilde{b}(f)}{S(f)} e^{-i2\pi ft}, \quad (3.36)$$

which allows us to define a *complex* SNR timeseries $z(t)$

$$z(t) = \langle s|h_T \rangle(t) = (s|h_R)(t) + i(s|h_I)(t) \quad (3.37)$$

so that the SNR timeseries has a compact expression: $\rho(t) = |z(t)|$.

Using the properties of the cross-correlation, we can write the complex SNR timeseries as the cross-correlation of the *whitened* $s^W(t)$ detector's output and the whitened templates h_R^W, h_I^W :

$$z(t) = (s^W \star h_R^W)(t) + i(s^W \star h_I^W)(t) \quad (3.38)$$

$$= \int_{-\infty}^{\infty} d\tau s^W(\tau) h_R^W(\tau - t) + i \int_{-\infty}^{\infty} d\tau s^W(\tau) h_I^W(\tau - t). \quad (3.39)$$

The expression is remarkable in its simplicity: the time dependent matched-filtering statistic is the quadrature sum of the cross-correlation of the whitened data with the two whitened normalized orthogonal templates.

Statistical distribution of $\rho(t)$ A matched-filtering search may employ Eq. (3.38) (or equivalently Eq. (3.35)) to compute the SNR timeseries for a large number of templates and it may claim a detection whenever any of the SNR timeseries exceeds a certain threshold value $\bar{\rho}^2$.⁸ In this case, a *trigger* is generated. Of course, for any choice of threshold $\bar{\rho}^2$, there is a small but non vanishing probability to obtain a SNR value above the threshold even in the absence of a signal, i.e. due to noise fluctuations. For this reason, it is important to characterize the statistical properties of $\rho(t)$ in the case of noise only: this will allow to measure the statistical significance of a trigger, by assigning to each a false alarm probability or, equivalently, a false alarm rate.

⁸As we will see below, the GstLAL pipeline does not impose a threshold on ρ^2 but rather on a ranking statistic \mathcal{L} , which takes into account other quantities besides the SNR value.

For a single real *normalized* template \hat{h} , in the presence of noise only ($s = n$), the random variable $(\hat{h}|s)$ can be seen as the sum of an “infinite” number of Gaussian random variables, hence the distribution of $(\hat{h}|s)$ also follows a Gaussian with 0 mean and unit norm:

$$(\hat{h}|s) \sim \mathcal{N}(0, 1) \quad (3.40)$$

where the variance, which in principle is arbitrary, is set by the filter normalization. The random variable $\rho^2 = |x + iy|^2$ is then the quadrature sum of two random normal variables $x, y \sim \mathcal{N}(0, 1)$ hence it follows a χ^2 distribution with two degrees of freedom $\rho^2 \sim \frac{1}{2}e^{-\rho^2/2}$, which can be written as a distribution for ρ through a change of variables

$$p(\rho) = \rho e^{-\rho^2/2}. \quad (3.41)$$

For each trigger with SNR $\bar{\rho}$, we can then define its *false alarm probability* (FAP) as the probability of obtaining a value of $\rho > \bar{\rho}$ only due to noise fluctuations

$$FAP(\bar{\rho}) = \int_{\bar{\rho}}^{\infty} d\rho \rho e^{-\rho^2/2}. \quad (3.42)$$

In Fig. 3.1 we show a simulated SNR timeseries for an injected signal, where the presence of a trigger is manifest by the large peak in the timeseries. We also display a histogram of the simulated SNR values. We note that the bulk of the distribution is formed by background triggers, generated by random noise fluctuations and whose distribution agrees with Eq. (3.41). In correspondence of the injected signal, the plot shows a number of large SNR outliers from the expected distribution, with a very low false alarm probability.

3.3 Template Banks

As discussed above, the search statistic can be maximized analytically only over some of the 17 parameters characterizing a generic eccentric BBH. The maximization with respect to the others is performed by a brute force approach, where the search statistic is evaluated on a discrete set of values. In practice, this means that the data are filtered multiple times with different templates and for each template, the SNR timeseries is analysed looking for peaks, corresponding to potential candidates. As already mentioned, the set of templates used to filter the data is called *template bank* [63, 325, 330–333]. In what follows, we will denote by θ the parameters for which a brute force approach is required, as opposed to the whole 17 parameters describing a BBH, which we collectively referred to as ϑ .

Of course, using a discrete set of templates, implies a reduced search sensitivity. Indeed, by filtering a signal, characterized by parameters θ_2 , with the “wrong” template, characterized by parameters θ_1 , the measured value of the search statistic is

only a fraction of the value obtained by filtering the data with the appropriate template θ_2 , thus reducing the statistical significance of any potential discovery. This is quantified by the *match* $\mathcal{M}(\theta_1, \theta_2)$, defined as the fraction of SNR measured when filtering a signal $\hat{h}_p(\theta_2)$ with the template $\hat{h}_p(\theta_1)$:

$$\mathcal{M}(\theta_1, \theta_2) = \max_t \left| \langle \hat{h}_p(\theta_1) | \hat{h}_p(\theta_2) \rangle(t) \right|^2 \quad (3.43)$$

$$= \max_t \left| 4 \int_{f_{\min}}^{f_{\max}} df \frac{\tilde{h}_p^*(f; \theta_1) \tilde{h}_p(f; \theta_2)}{S_n(f)} e^{i2\pi f t} \right|^2 \quad (3.44)$$

where we considered a frequency window $[f_{\min}, f_{\max}]$ and we considered *normalized* templates. In the template bank jargon, the quantity $\left| \langle \hat{h}_p(\theta_1) | \hat{h}_p(\theta_2) \rangle(t) \right|^2$ is also called *overlap*, which we denote by $\mathcal{O}(\theta_1, \theta_2, t)$. The match has values in $[0, 1]$, where is equal to one only when $\theta_1 = \theta_2$. In this definition, we assumed that the templates correspond to quadrupolar circular aligned-spin systems, hence the match is defined from the search statistic Eq. (3.25). In the next Chapter, we will drop this assumption and we will define a match based on the general statistic Eq. (3.30).

With the match at hand, we can define the *fitting factor* of a signal θ as the maximum match of the signal $\hat{h}_p(\theta)$ with the templates of the bank [334]:

$$FF(\theta) = \max_{\theta' \in \text{bank}} \mathcal{M}(\theta, \theta'). \quad (3.45)$$

This quantity is useful for the purpose of bank validation, as it gives important pieces of information about the performance of the bank in different regions of the parameter space. Borrowing the jargon of GW searches, we call *injections* the simulations for which we evaluate the fitting factor. Moreover, we can use the match to define a distance d between two points θ_1, θ_2 of the parameter space:

$$d(\theta_1, \theta_2) = \sqrt{1 - \mathcal{M}(\theta_1, \theta_2)}. \quad (3.46)$$

As discussed below with more details, many template placement algorithms rely on this definition for optimal placement.

The generation of a template bank requires to find a delicate compromise between two competing needs. On the one hand, it is desirable to have a fine template spacing, so that the SNR loss of a random BBH signal due to the discreteness of the bank is low: this would require to place as many templates as possible. On the other hand, a high number of template results in a high computational cost and, more importantly, in a larger number of degrees of freedom to search over, resulting in more matches to false alarms, which could possibly reduce the search sensitivity. Thus, an optimal template bank is composed of the smallest number of templates that guarantees that

only a small fraction of SNR from GW signals is missed due to the discreteness of the template bank [335]. The maximum tolerable SNR loss is set by the *minimal match* (MM) parameter. The minimal match is approximately related to the fraction of astrophysical signals lost due to the discreteness of the bank: assuming that BBHs are uniformly distributed on the volume and since the SNR scales inversely with the binary distance, the rate of lost events roughly scales as FF^3 [331]. It is common to set $MM = 0.97$, so that less than 10% of the signals are expected to be missed because of the template bank.

Once a minimal match is chosen, it is crucial to choose the templates in such a way that they cover the space of interest “as uniformly as possible”: this makes sure that the required bank accuracy is achieved with the lowest number of templates possible, hence at a low computational cost. A number of methods have been developed to effectively place the templates: their effectiveness and speed heavily depends on the dimensionality of the template bank, i.e. on the dimension of the parameter vector θ , and on the waveform model being used. Below, we provide a succinct review of the most important methods.

Metric template placement Historically, the first template placement algorithms was developed to place templates on a lattice with constant spacing [331,336,337] on a suitable manifold. The geometry of the manifold is described by a scalar product between templates, represented by a metric $M_{ij}(\theta)$:

$$M_{ij}(\theta) = -\frac{1}{2} \left. \frac{\partial^2 \mathcal{M}(\theta, \theta + \Delta\theta)}{\partial_i \Delta\theta \partial_j \Delta\theta} \right|_{\Delta\theta=0}. \quad (3.47)$$

The metric above is designed to provide a first order approximation to the distance Eq. (3.46), so that:

$$d^2(\theta, \theta + \Delta\theta) = 1 - \mathcal{M}(\theta, \theta + \Delta\theta) \simeq M_{ij}(\theta) \Delta\theta_i \Delta\theta_j. \quad (3.48)$$

For simple analytical PN waveform models, the expression for the metric is known analytically [336], provided that also the PSD admits an analytical expression. This makes the metric evaluation particularly convenient. Thus, the problem of template placement is cast into the problem of covering a non-flat manifold, such that every point of the manifold is closer than $\sqrt{1 - MM}$ to at least one lattice point. In other words, every point in parameter space has to be closer than $\sqrt{1 - MM}$ to at least one template – it has to be *covered* by at least one template.

While several template placement schemes have been developed to place templates, all differing in the details, they are mostly based on two key ideas:

1. Finding a coordinate system were the metric is approximately flat $M_{ij}(\theta) \simeq \delta_{ij}$. This allows to use the large variety of lattice geometries available for the eu-

clidean space [338]. To this goal, the two BH masses m_1, m_2 are often parameterized using pairs of chirp times variables, such as τ_0, τ_1 [331] or τ_0, τ_3 [339].

2. Finding a suitable lattice to cover the parameter space, with several lattice being employed from a simple cubic lattice [335,336] to a more complicated geometries [339–343].

Despite being appealing for its simplicity and optimality under ideal conditions, the geometric placement suffers from sever drawbacks, limiting its applicability. First of all, obtaining coordinate transformation which avoids varying metric components has proved particularly difficult, especially for modern approximants, which include sophisticated effects. Secondly, and perhaps more importantly, the geometric template placement suffers from poor scaling with the number of dimensions. Originally developed for searches of non-spinning system covering only a two dimensional space, producing a lattice in high dimensions is extremely costly. Moreover, as pointed out in [344], for a large number of dimension, the best performing lattices show only marginal improvement over a simple cubic lattice.

Random template placement An appealing alternative to geometric placement relies on generating random template bank [345]. While the first method aims at covering every point in the relevant parameter space, the latter only covers each point with a probability $\eta < 1$. Templates are randomly drawn from the uniform probability distribution induced by the metric, without checking for their mutual distance, until the required coverage fraction has been achieved [346]. This approach allows to greatly reduce the number of templates, as compared to the geometric approach, and at the same time it has a very good scaling with the number of dimensions. While the method may seem sub-optimal with respect to a lattice, in [345,347] it is argued that for high-dimensional spaces, random template banks outperform even the best known lattice in terms of coverage (at a fixed number of templates), effectively beating the “curse of dimensionality”.

Stochastic template placement While the random algorithm avoids many of the problems introduced by the geometric placement, it has a fundamental drawback, which has so far limited its applicability: as it does not check between the distance between templates, it covers the space with more templates than might be required with a more efficient sampling scheme. This forms the core idea of the stochastic template placement method [348,349]. Starting from an empty template bank, a template *proposal* is randomly drawn from a suitable distribution. The proposal is added to the template bank only if its fitting factor Eq. (3.45) with the templates already accepted is lower than the minimal match requirement MM . The iteration terminates after a

large number of proposals are consecutively rejected. The fitting factor can be computed either with the match Eq. (3.43) or with its metric approximation Eq. (3.48), with the latter being generally faster but less precise [350]. The stochastic method produces very high quality template banks and due to its simplicity and efficacy, it has become the standard tool for template bank generation [351] in mostly all the modern template banks used for GW searches [327, 328, 352–354].

Even though the stochastic approach has proven very powerful, its performance does not scale well with (i) the number of templates and, most importantly, (ii) with the number of dimensions of the parameter space. Handling a large number of templates can have a large impact on computing time and memory, because for every new proposal, a waveform needs to be generated and stored and many expensive match calculations need to be performed - modern template bank with millions of templates may take up to weeks to be generated. Furthermore, the sheer number of dimensions can have an even more catastrophic impact on the bank generation cost. Indeed, at every iteration the stochastic algorithm computes the distance between $N_p \sim r^D$ pairs of templates within a given radius r . It is clear how the number of match computations diverges for large dimensional spaces. Moreover, in a higher number of dimensions, the advantages of controlling the mutual distance between templates become less prominent, as demonstrated in [347] and also observed in our work presented in Ch. 4 and in [1]. In this scenario, as previously discussed, a random template bank approaches optimality. Therefore, the distance control executed by the stochastic algorithm may result in an unnecessary expenditure of computational resources. As a template bank for precessing systems necessarily gathers a large number of templates in a high number of dimension, the generation of a precessing template bank with the stochastic approach might prove unfeasible due to a huge computational cost. In the next chapter we show how we can efficiently generate a template bank of precessing signals by coupling the random template placement method with an effective parameter space sampling scheme.

3.4 Features of a Realistic Pipeline

The matched-filtering technique for aligned-spin circular binaries described above is implemented in the state-of-the-art matched-filtering pipelines which are routinely deployed for the search of LIGO, Virgo and KAGRA data. Notable examples are Gst-LAL [77–82], PyCBC [355–359], the Multi-Band Template Analysis (MBTA) [360, 361], the Summed Parallel Infinite Impulse Response (SPIIR) pipeline [362–364] among others [63, 365–369]. While they differ in many implementation details, they all utilize the matched filtering to achieve the same goal of producing a list of *events*. In

this context, an event is a flag for a segment of data that potentially identifies a GW signal. The events identified by each pipeline are then ranked from the most likely to the least likely to be a GW signal according to their probability of being recorded due to a noise fluctuation only.

A matched-filtering search can happen either *online* or *offline*. For an online configuration [81], the events are generated in a streaming fashion as soon as the data are available: while the pipeline sensitivity might be sub-optimal due to a poor background⁹ estimation, it allows to discover potential candidates with seconds of delays, thus enabling multi-messenger applications. On the other hand, the offline configuration is not concerned with low latency and it uses all the data available for a more thorough background estimation. It is usually run to confirm the results of an online search or to search for systems which are not expected to produce multi-messenger signatures, such as the intermediate mass black holes (IMBHs) [370–373] or sub-solar mass (SSM) black holes [374–378]. In what follows we will describe in details the *offline* GstLAL pipeline, used for our current work, highlighting the differences and similarities with other pipelines for matched-filtering GW searches.

3.4.1 Generating Triggers

The first step of the GstLAL pipeline is of course the filtering of the data. The pipeline computes the correlation between each template and data in time domain, using Eq. (3.38). While in frequency domain the cost of filtering D data-points with N templates is $ND \log D$, in time domain the filtering is more costly, being ND^2 . To reduce the computational cost associated to filtering, the GstLAL pipeline employs the LLOID method [77, 379], which consists in splitting the template bank in small groups of $O(500)$ templates and to perform the singular value decomposition (SVD) of the template $N \times D$ matrix H , whose rows correspond to different templates. The SVD provides a low rank approximation of H , where each template is written as a linear combination of K SVD basis:

$$H = \begin{bmatrix} \mathbf{h}_1 \\ \vdots \\ \mathbf{h}_N \end{bmatrix} = VU \quad (3.49)$$

where $\mathbf{h}_i \in \mathbb{R}^D$ is the template evaluated on a discrete time grid and V and U are two matrices with dimensions $N \times K$ and $K \times D$ respectively. If K is much smaller than the number of templates N , the computational cost of filtering is greatly reduced. The LLOID method also employs a time slice mechanism for the templates, so that

⁹Here we refer to the characterization of the triggers due to noise fluctuation and not associated to any astrophysical signal. As we will discuss in Sec. 3.4.3, these are also called “background” triggers.

the low frequency part of the templates is evaluated at a lower sample rate, with consequent computational benefits. Traditionally, the splitting of the bank into many different groups is done by tiling the $\mathcal{M}_c - \chi_{\text{eff}}$ space with rectangles, also called *SVD bins*. More recently, [380] introduced a splitting based on the PN variables μ^1 and μ^2 [381]. We will investigate this matter for precessing signals in Ch. 6.

As modern template banks consist of millions of templates, it is unfeasible to store the SNR timeseries with a typical sample rate of 4096 Hz for every template and, for this reason, the SNR timeseries is discretized into *triggers*. A trigger records the peak of the SNR timeseries within a 1 s window and it is only stored if the peak exceeds a threshold of $SNR = 4$. A trigger is a vector $\{\vec{t}, \vec{\Phi}_{\text{T}}, \vec{\rho}, \vec{\xi}^2, \vec{D}, \theta\}$ consisting of the time t of the recorded peak, the SNR ρ , the phase Φ_{T} of the complex SNR peak $z(t = t_0)$, the parameters θ of the matching template, the value of the signal consistency test ξ^2 (see next section) and the instantaneous sensitivity of each detector, measured in terms of horizon distance D . The latter is commonly defined to be the distance at which a $1.4 M_{\odot} - 1.4 M_{\odot}$ BNS is observed with a SNR equal to 8 [63].

The trigger generation phase is common to most of the pipelines, with some important differences. First of all, the vast majority of them computes the SNR time-series in frequency domain with Eq. (3.35), thus avoiding the SVD decomposition of the templates. Because of this, a different signal consistency test is employed, consisting in the traditional χ^2 test. Finally, different pipelines may use different time windows and thresholds for trigger generation.

It often happens that more than one detector is operating at every given time. In that situation, the data from every instrument are scanned with matched filtering and triggers are produced for every instrument. As a loud astrophysical signal is expected to produce a trigger in all the observing instruments, it is convenient to group together triggers from different instruments happening around the same time and from the same template and to assign them to the same event. This procedure produces a set *coincidences* [63], which we denote as a set of vectors

$$\{\vec{t}, \vec{\Phi}_{\text{T}}, \vec{\rho}, \vec{\xi}^2, \vec{D}, \theta\},$$

where each row of the vector refers to a different instrument and a coincidence can be formed even by a single trigger. Considering coincidences is a very powerful way to suppress the background of noise triggers and reduce the number of false positives returned by the search. Once coincidences are formed, the rest of the pipeline is devoted to rank the coincidences and to evaluate their significance: coincidences with high significance are eventually considered events.

3.4.2 Signal Consistency Test

A fluctuation of Gaussian noise or the presence of loud transient burst of noise can occasionally mimic the effect of a GW signal in the detector, leading to a high value of measured SNR. Clearly, this can produce a large number of false positives in the candidate list or may reduce the search sensitivity due to background contamination. Therefore, it is important to record an additional quantity (or statistic) that helps to distinguish between peaks in SNR due to astrophysical signals and to loud transient noise bursts [382]. To construct this quantity, we utilize the idea that, although a terrestrial noise burst may produce a high SNR, its time evolution in the SNR time series differs significantly from that of an expected astrophysical signal.

This idea is implemented within GstLAL through the introduction of the ξ^2 test [77], where the expected complex SNR timeseries $R(t)$ is compared with the measured SNR timeseries $z(t)$:

$$\xi^2 = \frac{\int_{-\delta t}^{\delta t} dt |z(t) - R(t)|^2}{\int_{-\delta t}^{\delta t} dt \left(2 - 2 \left|\frac{R(t)}{R(0)}\right|^2\right)} \quad (3.50)$$

where the integration extends in a small time window $[-\delta t, \delta t]$ around the trigger time. For circular aligned-spins systems without HMs, the expected SNR timeseries is readily computed by plugging a signal of the form of Eq. (3.23) into the search statistic Eq. (3.35):

$$R(t) = \frac{1}{D_{\text{eff}}} \langle e^{i\phi} h_p | \hat{h}_p \rangle(t) = \sqrt{(h_p | h_p)} \frac{e^{-i\phi}}{D_{\text{eff}}} \langle \hat{h}_p | \hat{h}_p \rangle(t) \quad (3.51)$$

$$= z(0) \langle \hat{h}_p | \hat{h}_p \rangle(t) \quad (3.52)$$

where in the last line we recognized that $z(0) = \sqrt{(h_p | h_p)} \frac{e^{-i\phi}}{D_{\text{eff}}}$. The quantity $\langle \hat{h}_p | \hat{h}_p \rangle(t)$ is also called *template autocorrelation*. The value of ξ^2 is recorded for every trigger and it is used for the subsequent ranking process. The expression Eq. (3.52) for the expected SNR timeseries is not valid in the precessing and/or HM case and in Ch. 5 we will compute an expression valid in the general case.

Of course, other tests are available in the literature [366, 382–387] and they all have been employed in real searches. An example of widely utilized test is the χ^2 signal consistency test [382], which splits the template into p frequency bins which contributes for equal power to the total SNR ρ . The test then computes the χ^2 statistic, which checks whether the measured power ρ_i^2 in each frequency bin is equal to the expected power ρ^2/p :

$$\chi^2 = p \sum_{i=1}^p \left[\left(\frac{\rho_p^2}{p} - \rho_{p,i}^2 \right)^2 + \left(\frac{\rho_c^2}{p} - \rho_{c,i}^2 \right)^2 \right] \quad (3.53)$$

where $\rho_p^2 = (s|\hat{h}_p)^2$ and similarly for ρ_c^2 . The test is employed by the PyCBC pipeline (among others) to reweight the matched filter SNR, hence to downrank the significance of triggers with high values of χ^2 .

3.4.3 Ranking

In modern pipelines, hundreds of thousands of triggers are generated every second. As the vast majority of such triggers arise from noise fluctuations, an effective search must be able to distinguish the triggers of astrophysical origin (signal events), from the *background* of coincidences originated from noise (noise events). To do so, every search develops a *ranking statistic* \mathcal{L} , which is a numerical value assigned to each coincidence. A good ranking statistic assigns a large value to signal events and a low value to noise events: in this way, the ranking statistic can be used to order coincidences from the most likely to the least likely to be of astrophysical origin. Moreover, a threshold on \mathcal{L} can be used to claim a detection. In principle, one could use the search statistic ρ as a ranking statistic. However, ρ is the optimal statistic only in Gaussian noise and in the presence of non-Gaussianities or loud transient noise burst, we need a more powerful statistic, which takes into account all the other information available.

Unlike other pipelines, the GstLAL pipeline builds a ranking statistic by computing the logarithmic ratio between the likelihood $p(\dots | \text{signal})$ that a coincidence is the result of a signal and the probability $p(\dots | \text{noise})$ that each coincidence is a result of a random noise fluctuations [79, 82, 388, 389]:

$$\mathcal{L} = \log \frac{p(\vec{t}, \vec{\Phi}_T, \vec{\rho}, \vec{\xi}^2, \vec{D}, \theta | \text{signal})}{p(\vec{t}, \vec{\Phi}_T, \vec{\rho}, \vec{\xi}^2, \vec{D}, \theta | \text{noise})} \quad (3.54)$$

where ρ is a shorthand for SNR and $\{\vec{D}, \vec{\rho}, \vec{\xi}^2, \vec{t}, \vec{\Phi}_T\}$ defines a coincidence of one or more instruments. Following the GstLAL jargon, we call *signal model* the statistical model for $p(\dots | \text{signal})$, while *noise model* refers to $p(\dots | \text{noise})$. Note the that both signals and noise models depend on the template parameter θ . We may also refer to \mathcal{L} as the likelihood ratio (LR).

Both the signal and noise model are parametric probabilistic models, whose parameters are set at the time of the search. They are conveniently factorized to exploit the known correlations between the recorded quantities — see also [82, Eqs. (2) and (8)]. While in principle the LR may depend on the template parameter θ , it is more convenient to divide the template bank into *background bins* and have the LR depend only on the background bin of a given template. This choice reduces the number of free parameters in the LR model and it is found to improve its performance [389]. Typically, one background bin corresponds to one or more SVD bins. The robustness

of the ranking statistic can be enhanced by considering pieces of information regarding the status of the detector, hence downranking triggers originated at times where the detector has known data quality issues [390, 391].

Once the ranking statistic \mathcal{L} has been computed, the coincidences are clustered within a 4 s time window and only the one with maximum likelihood is retained. This makes sure to remove multiple triggers from the same astrophysical signal and it avoids to propose multiple candidates events for a single trigger time.

3.4.4 False Alarm Rate

To claim a detection, it is customary to put a threshold $\bar{\mathcal{L}}$ on the ranking statistic and to label as “signal” all the triggers with $\mathcal{L} > \bar{\mathcal{L}}$. As also discussed in Sec. 3.2.3, there is always a small but non vanishing probability that a high value of \mathcal{L} is recorded due to only noise fluctuations. For this reason, claiming a detection with a threshold on \mathcal{L} must be accompanied by an estimation of the probability that the claimed “signal” is actually of astrophysical origin. In other words, we need an estimate of the probability of a false claim or, equivalently, we need a measure of the statistical significance of a detection.

Similarly to what done in Sec. 3.2.3 for triggers in Gaussian noise, we can measure the statistical significance of an event with likelihood $\bar{\mathcal{L}}$ by comparing the event with the distribution $p(\mathcal{L}|\text{noise})$ of the LR values for noise triggers. This allows us to compute the *false alarm probability* (FAP) of recorded event, defined as the probability that a search targeting only noise yields a trigger with a \mathcal{L} greater than the recorded trigger $\mathcal{L} > \bar{\mathcal{L}}$. Traditionally, $p(\mathcal{L}|\text{noise})$ is estimated using the time slide method [365], which consist in producing artificial coincidences between time shifted triggers and evaluating their ranking statistic. The GstLAL pipeline however estimates $p(\mathcal{L}|\text{noise})$ by integrating the noise likelihood $p(\dots|\text{noise})$ on the surfaces $\Sigma(\mathcal{L})$ of constant \mathcal{L} [388, 389]:

$$p(\mathcal{L}|\text{noise}) = \int_{\Sigma(\mathcal{L})} d^{n-1}\vec{\Theta} p(\vec{\Theta}|\text{noise}) \quad (3.55)$$

where $\vec{\Theta}$ is a shorthand for all the quantities defining a trigger \vec{t} , $\vec{\Phi}_T$, $\vec{\rho}$, $\vec{\xi}^2$, \vec{D} , θ and $\Sigma(\mathcal{L})$ is mathematically defined as

$$\Sigma(\mathcal{L}) = \left\{ \vec{\Theta} \text{ s.t. } \log \frac{p(\vec{\Theta}|\text{signal})}{p(\vec{\Theta}|\text{noise})} = \mathcal{L} \right\}. \quad (3.56)$$

The procedure of estimating $p(\mathcal{L}|\text{noise})$ is crucial for every GW search and it goes under the name of *background estimation*. Once the background of the search $p(\mathcal{L}|\text{noise})$ is known, we can compute the FAP of a trigger with ranking statistic \mathcal{L} as:

$$FAP(\mathcal{L}) = \int_{\mathcal{L}}^{\infty} d\mathcal{L}' p(\mathcal{L}'|\text{noise}). \quad (3.57)$$

For a typical search, it is convenient to express the significance in terms of the *false alarm rate* (FAR), defined for any \mathcal{L} as the average frequency of occurrence of a noise event ranking statistic $> \mathcal{L}$:

$$FAR(\mathcal{L}) = \frac{FAP(\mathcal{L})}{T} = \frac{1}{T} \int_{\mathcal{L}}^{\infty} d\mathcal{L}' p(\mathcal{L}'|\text{noise}) \quad (3.58)$$

where T is the length in time of the search.

The need to introduce the FAR can be better understood in the context of the theory of hypothesis testing. To make a decision regarding the origin of a given trigger, it is natural to introduce the null hypothesis that “*the event with $\mathcal{L} = \bar{\mathcal{L}}$ has originated from a noise source*” and test whether the hypothesis is confirmed by the data through the computation of a p-value. A standard textbook result guarantees that in this scenario, the correct p-value to use is indeed the false alarm probability Eq. (3.57). Of course, claiming a detection always amounts to choosing a somewhat arbitrary threshold on some statistic, which in our case can be either the false positive rate or the ranking statistic¹⁰. A detection threshold on the FAP (or equivalently on the FAR) has the benefit of being more robust as it is based on the statistical properties of \mathcal{L} , which depend on the search performance and not on an arbitrary definition of the LR. A typical choice for a FAR threshold is 1/year.

3.4.5 p_{astro}

Another way to measure the statistical significance of an event, hence to make a more informed choice for the detection threshold, consist on defining for each event the probability of being of astrophysical origin p_{astro} . This approach is somehow complementary to computing the FAP, as the latter is intimately related to the probability that an event has originated from noise. First introduced in [294], the method to estimate p_{astro} relies on a Bayesian framework to estimate the rate of signal and noise events and use this information to infer the origin of a single event. It starts from the assumption a search produces both signal and noise events according to two independent Poisson processes. The density of the observed number of events N is then given by the superposition of the two processes [392, 393]:

$$\frac{1}{T} \frac{dN}{d\mathcal{L}} = R_n p(\mathcal{L}|\text{noise}) + R_s p(\mathcal{L}|\text{signal}) \quad (3.59)$$

where R_n, R_s are the rates of noise and signal triggers, T is the time of the search and $p(\mathcal{L}|\text{signal})$ can be estimated with Eq. (3.55) by replacing the noise model with the signal model.

¹⁰Of course, the two are equivalent since the FAP is an monotonous function of \mathcal{L} , so a threshold in FAR is directly translated into a threshold on \mathcal{L} .

Given an observed set of triggers $\{\mathcal{L}_i\}_{i=1}^N$ above a threshold $\bar{\mathcal{L}}$, chosen such that noise triggers dominates at the threshold, we can write down a posterior distribution for the trigger rates R_n, R_s [294, 394]:

$$p(R_n, R_s | \{\mathcal{L}_i\}_{i=1}^N) \propto \left\{ \prod_{i=1}^N [R_n p(\mathcal{L}_i | \text{noise}) + R_s p(\mathcal{L}_i | \text{signal})] \right\} \times \exp [-T (R_s + R_n)] \frac{1}{\sqrt{R_n R_s}} \quad (3.60)$$

where we used the Jeffreys prior $p(R_n, R_s) \propto (R_n R_s)^{-1/2}$ for the likelihood in consideration. An estimation from the rates can be obtained by employing stochastic sampling method available in the literature.

With the posterior at hand, we can define a the conditional probability that an event with detection statistic \mathcal{L} is of astrophysical origin p_{astro} [294, 395]:

$$p_{\text{astro}} = \frac{\bar{R}_s p(\mathcal{L} | \text{signal})}{\bar{R}_n p(\mathcal{L} | \text{noise}) + \bar{R}_s p(\mathcal{L} | \text{signal})} \quad (3.61)$$

where the $\bar{R}_{s/n}$ denotes the posterior averaged rate of signal and noise triggers respectively [396]. The same computation can be generalized to multiple classes [294], to tackle the case the experimenter is interested in distinguishing between multiple sources of astrophysical signals, such as a BBH, a NSBH and BNS.

Due to its simplicity and interpretability, it is common to claim a detection based on p_{astro} rather than on the FAR. Indeed, modern catalogues [16, 18] produced by the LIGO-Virgo-KAGRA collaboration are constructed by including all the events with $p_{\text{astro}} > 0.5$. Despite this, the FAR still remain an important quantity to monitor before claiming a detection, mostly because it does not depend on any assumption about the nature of signal being detected.

CHAPTER 4

Generating Template Banks in High Dimensional Spaces

4.1 Introduction

As gravitational-wave astronomy enters a mature state with an ever increasing instrument sensitivity, the accessible parameter space of BBH mergers continues to grow. Besides standard aligned-spin GW searches for stellar-mass BBH mergers [15–18], there are GW searches targeting the parameter space of sub-solar mass black holes (SSM) [376,397,398], primordial BHs [375], eccentric binaries [68,166,375,399–403] and intermediate-mass BHs (IMBH) [370,371,373]. Moreover, there is a growing interest in GW searches for more complex binaries, such as those with precession [66,74,185,327,328,404–406] or higher-order mode (HMs) content [328,372,407–409].

As discussed in the previous chapter, the most sensitive GW searches for signals from compact binary mergers relies on filtering the data with a template bank of model waveforms. The task of generating a template bank is far from trivial, as one must balance good coverage (thus ensuring low SNR loss due to the discreteness of the bank) with maintaining a manageable bank size. We saw that this is very effectively achieved by the *stochastic* method [348,349], which iteratively adds templates to the bank by checking that they are not “too close” to each other. However, the stochastic approach becomes computationally and memory expensive as the bank size and the dimensionality of templates grow. As the BBH searches grow in complexity due to the inclusion of more physical effects and hence more dimensions, the

stochastic approach struggles to produce template banks in a feasible amount of time and with a manageable memory footprint. This poses the challenge of finding a viable alternative for template bank generation, which is able to deliver large banks in a high-dimensional parameter space, such as those associated with precession, eccentricity and HMs.

Revitalizing a pioneering line of research [345], the *random template placement* is appealing to cover such high dimensional spaces due to its speed and good scaling properties, as demonstrated in [347]. The strength of the method is twofold: on the one hand, since no distance between templates is computed, the template placement is very fast and memory efficient; on the other hand, by only covering a fraction of the space, the number of templates remains under control. Moreover, the cost does not increase for an increasing number of dimensions. Generating a random template banks requires the ability to effectively sample templates “uniformly” across the parameter space. Traditionally, due to the high dimensionality of the space, expensive sampling techniques, such as Markov chain Monte Carlo, must be used. This poses a serious limitation to the range of applicability of the method — and perhaps limited its application in GW data analysis so far. Without a fast sampling method, the speed up promised by the new method is washed away by the cost of a large number of metric evaluations.

In this chapter based on [1], we address the challenges described above by covering high-dimensional spaces with random template banks. As a first step, we derive a novel expression for the metric, which is suitable for generic precessing and/or HM waveforms. In doing so, we drop several symmetry assumptions that enters the standard metric computation. The metric is then expressed in terms of the gradients of the waveform. Secondly, to enable a fast template sampling, we employ machine learning and train a *normalizing flow* model to efficiently sample templates from the parameter space. While the first innovation delivers an accurate distribution for the templates throughout the space, the use of a normalizing flow allows us to generate random template banks in a few hours (including the training time).

The combination of a new metric expression and the normalizing flow model, applied to the random template placement algorithm, makes our method particularly well-suited for dealing with high-dimensional ($> 4D$) parameter spaces, such as those associated with precessing or eccentric searches. Our method is implemented in an open-source, production-ready, Python package `mbank` [410], available on GitHub¹ and on the PyPI repository².

The rest of the chapter is devoted to the presentation and description of our meth-

¹<https://github.com/stefanoschmidt1995/mbank>.

²The package is distributed under the name `gw-mbank`.

ods and software package. In Sec. 4.2 we present the details of our bank generation algorithm. In Sec. 4.3 we assess the accuracy of our template placing method in all its parts. Furthermore, we reproduce two banks available in the literature [328, 380] created with independent codes: this will be the topic of Sec. 4.4. To demonstrate the capabilities of `mbank`, in Sec. 4.5, we present two large banks covering “exotic” regions of parameter space: a precessing bank and an IMBH bank with HM content. We also discuss some possible further applications of our normalizing flow model, including a study of the size of the precessing neutron star-black hole (NSBH) parameter space. Finally, in Sec. 4.6 we discuss some possible future development of our work and gather some final remarks in Sec. 4.7.

Throughout the chapter we will use the term “standard” to refer to the searches for circularized, aligned-spin BBHs without imprints of HMs, currently conducted by the LIGO-Virgo-KAGRA collaboration.

4.2 Methods

A matched-filtering search aims at maximizing the search statistic over a large number of templates, where each template depends on a number of parameters describing the properties of the BBH being searched. We remind the reader that in the general case, the search statistic Eq. (3.30) depends on 12 parameters: they are the two BH masses m_1, m_2 , the two three-dimensional spins $\mathbf{s}_1, \mathbf{s}_2$, the inclination angle ι , the reference phase ϕ and the eccentricity parameters e, a . On the other hand, if we restrict our attention to circular aligned-spin BBH without HM, it is sufficient to consider the search statistic Eq. (3.26), which only depends on four parameters.

For the purpose of template placement, it is useful to think of the parameter space of BBH signals as a D -dimensional manifold \mathcal{B}_D , embedded in a large 12 dimensional manifold \mathcal{B} . Each point of the manifold corresponds to a GW signal. The number of dimensions D depends on the BBH variables under consideration. As the parameters that do not enter the interesting space can be freely neglected (i.e. set to 0 or to a meaningful constant value), the manifold \mathcal{B}_D is effectively a lower dimensional *projection* of the large manifold \mathcal{B} .

To place templates on \mathcal{B}_D , it is standard to equip the manifold with a distance (called *mismatch*), which also naturally defines a volume element at every point in space. The volume element defines the “uniform” probability distribution according to the metric. A random template bank will be populated by templates drawn from such distribution, until a certain coverage is reached. For this reason, our primary concern is to sample from the manifold and to check for coverage. To effectively do so, we rely on the three steps below:

1. Construction of a metric approximation of the match between templates. This makes \mathcal{B}_D a Riemannian manifold with a volume element.
2. Training of a normalizing flow model to sample from the manifold.
3. Placing the templates by sampling from the normalizing flow model and checking for coverage, following [346].

The rest of this section details the steps above.

4.2.1 The Metric

The definition of a metric on the manifold \mathcal{B}_D provides a fast-to-compute approximation to the *mismatch* (distance) between templates and an estimation of the volume element at each point in the space. While all the literature available [335, 336, 338, 341, 343, 345, 346] is built only for the “standard” case Eq. (3.26), in this work we will construct a metric approximation of the “symphony” search statistic Eq. (3.30): to our knowledge, this is done here for the first time. Clearly, the metric expression we derive reduces to the standard case when $(\hat{h}_+|\hat{h}_\times) = 0$.

Before introducing the metric, we need to generalize the definition of the match Eq. (3.43) to the general case. Given two points of the manifold θ_1, θ_2 , we begin by defining the “symphony” overlap $\mathcal{O}_{\text{sym}}(\theta_1, \theta_2, t)$ between normalized precessing and/or HM templates as:

$$\mathcal{O}_{\text{sym}}(\theta_1, \theta_2, t) = \frac{1}{1 - \hat{h}_{+\times}(\theta_2)^2} \left\{ \left(\hat{h}_+(\theta_1) e^{i 2\pi f t} |\hat{h}_+(\theta_2) \right)^2 + \left(\hat{h}_+(\theta_1) e^{i 2\pi f t} |\hat{h}_\times(\theta_2) \right)^2 \right. \\ \left. - 2 \hat{h}_{+\times}(\theta_2) \left(\hat{h}_+(\theta_1) e^{i 2\pi f t} |\hat{h}_\times(\theta_2) \right) \left(\hat{h}_+(\theta_1) e^{i 2\pi f t} |\hat{h}_+(\theta_2) \right) \right\} \quad (4.1)$$

where $\hat{h}_+(\theta) e^{i 2\pi f t}$ is the plus polarization $\hat{h}_+(\theta)$ translated by a constant time shift t and $\hat{h}_{+\times}(\theta) = (\hat{h}_+(\theta) |\hat{h}_\times(\theta))$. The “symphony” overlap retains the same physical interpretation of the standard case and amounts to the fraction of SNR recovered when filtering a signal $s = h_+(\theta_1)$ with a template evaluated at a point θ_2 , as a function of a time shift t .

In Eq. (4.1), we choose to compare the plus polarization of the first template with both polarizations of the second template. We are forced to make such arbitrary choice by the fact that the search statistic in Eq. (3.30) does depend on F_+, F_\times , unlike the standard case Eq. (3.25). Thus, if we don’t want the overlap to depend on two arbitrary combination coefficients, an arbitrary choice for the signal s is needed. Of course, any linear combination of $h_+(\theta_1)$ and $h_\times(\theta_1)$ works but we set $s = h_+(\theta_1)$ for computational convenience. Some numerical studies that we performed as part

of our development efforts (but not published) show that replacing $h_+(\theta_1)$ with any linear combination does not have a large impact on the metric definition below.

In the case of a “standard” search, $\hat{h}_{+\times} = 0$ and $\tilde{h}_+ = i\tilde{h}_{\times}$, hence the “symphony” overlap simplifies to:

$$\mathcal{O}_{\text{sym}}(\theta_1, \theta_2, t) = \mathcal{O}(\theta_1, \theta_2, t) = \left| \langle \hat{h}_+(\theta_1) e^{i 2\pi f t} | \hat{h}_+(\theta_2) \rangle \right|^2. \quad (4.2)$$

and agrees with the overlap we introduced in the discussion following Eq. (3.43). Note that, since Eq. (3.25) is symmetric³ between signal and template, the expression for the overlap in the “standard” case is also symmetric. This means that an arbitrary choice on the signal composition is no longer needed, as was the case for Eq. (4.1).

Relying on the expression for the overlap introduced above, we obtain a straightforward generalization to the precessing and/or HM case of the *match* $\mathcal{M}(\theta_1, \theta_2)$ introduced in Eq. (3.43). Therefore, we define the “symphony” match $\mathcal{M}_{\text{sym}}(\theta_1, \theta_2)$ between two templates evaluated at different points θ_1, θ_2 of the manifold as:

$$\mathcal{M}_{\text{sym}}(\theta_1, \theta_2) = \max_t \mathcal{O}_{\text{sym}}(\theta_1, \theta_2, t). \quad (4.3)$$

Such definition of the match between template, which combines Eq (4.3) with Eq. (4.1), is introduced in our work [1] for the first time. As it is directly tied to the relevant search statistic, using the “symphony” match for bank generation yields precessing and/or HM template banks of higher quality compared to those produced by the “standard” match definition. For this reason, the “symphony” match will be consistently used for the remaining of our work, and for simplicity we will refer to it simply as the match $\mathcal{M}(\theta_1, \theta_2)$.

Even though in general the match is not symmetric and does not satisfy triangular inequality, we can use it to introduce a *distance* d between two points on the D-manifold \mathcal{B}_D

$$d^2(\theta_1, \theta_2) := 1 - \mathcal{M}(\theta_1, \theta_2). \quad (4.4)$$

which can be approximated locally by a bilinear form d_M

$$d_M^2(\theta_1, \theta_2) := M_{ij}(\theta) \Delta\theta_i \Delta\theta_j \simeq 1 - \mathcal{M}(\theta_1, \theta_2). \quad (4.5)$$

The bilinear form d_M is represented by a D-dimensional square matrix $M_{ij}(\theta)$, defined at each point of the manifold, which we call *metric*. By comparing the two equations above, we identify the metric with the quadratic term of the Taylor expansion of $d_M(\theta + \Delta\theta, \theta)$ around $\Delta\theta \simeq 0$:

$$M_{ij}(\theta) = -\frac{1}{2} \left(H_{ij} - \frac{H_{ti} H_{tj}}{H_{tt}} \right) \quad (4.6)$$

³Indeed, for a “standard” signal $s \propto h_+$, hence $\hat{s} = \hat{h}_+$, and Eq. (3.25) does not depend on the antenna patterns functions, if s is normalized.

where $H(\theta)$ is the Hessian [331] of the overlap in Eq. (4.1), which is a $D + 1$ square matrix whose entries are explicitly given by

$$H_{ij}(\theta) = \frac{\partial_{ij}\mathcal{O}(\theta, \theta + \Delta\theta)}{\partial_i \Delta\theta \partial_j \Delta\theta} \Big|_{\Delta\theta=0}. \quad (4.7)$$

See App. 4.A for more details on the Hessian computation. Note that the metric is positive definite (i.e. has positive eigenvalues), since it is evaluated at the local minimum of a differentiable function.

A convenient expression for H in terms of the gradients of the waveform is presented in App. 4.A, with the full expression given in Eqs. (4.35)-(4.37). While identifying the metric with the Hessian is well motivated and yields reliable results, other definitions for M_{ij} are possible; this is briefly discussed in App. 4.B.

For most of the waveform models available, the gradients can be evaluated with finite difference methods. For a limited number of machine-learning based models [5, 6, 268, 270, 411], the gradients are available analytically.

Equipped with the metric from Eq. (4.6), the manifold \mathcal{B}_D becomes a Riemannian manifold with line element:

$$ds^2 = M_{ij}(\theta) d\theta_i d\theta_j. \quad (4.8)$$

We can then use standard results from differential geometry to compute distances and volumes. In particular, the volume of a subset \mathcal{T} of the manifold can be computed as:

$$\text{Vol}(\mathcal{T}) = \int_{\mathcal{T}} d^D\theta \sqrt{\det M(\theta)}, \quad (4.9)$$

where $\det M(\theta)$ is the determinant of the matrix $M_{ij}(\theta)$, also denoted as $|M|$. Moreover, we introduce the uniform probability measure, such that $p(V) \propto \text{Vol}(V)$ for any $V \subseteq \mathcal{B}_D$. The measure has the following probability distribution function (PDF):

$$p(\theta) \propto \sqrt{\det M(\theta)}. \quad (4.10)$$

Samples from the uniform distribution tend to have a “uniform” (i.e. constant) spacing, computed with the metric distance. Owing to this feature, the uniform distribution is a natural candidate to draw templates from.

4.2.2 Sampling from the Manifold

To generate a random template bank, we need to sample points on the manifold \mathcal{B}_D from Eq. (4.10). A simple way to do so is by means of a Markov Chain Monte Carlo (MCMC). However, this turns out to be unfeasibly expensive, since to obtain a single

sample, the metric must be evaluated tens of times. For instance, to produce a bank with $\mathcal{O}(10^6)$ templates, $\mathcal{O}(10^7)$ metric evaluations are required.

To speed up the sampling, we introduce a *normalizing flow* model. As we will show below, in order to train the model $\mathcal{O}(10^5)$ metric evaluations are sufficient: this is a small fraction of the metric evaluations needed to run a MCMC. Once trained, the normalizing flow model produces high quality samples from Eq. (4.10) in a small amount of time, effectively providing templates to populate a random template bank.

A normalizing flow model [412–415] is a machine learning model widely used to reproduce and/or parameterize complicated probability distributions. Mathematically, a flow is an *invertible* parametric function Φ_W which is trained to map samples θ from an arbitrary probability distribution $p(\theta)$ to samples \mathbf{x} from a multivariate standard normal distribution $\mathcal{N}(\mathbf{x}|0, \mathbf{1})$. The space of the \mathbf{x} is sometimes referred to as *latent space*. The parameters W of the flow are set in such a way that:

$$\mathbf{x} = \Phi_W(\theta) \sim \mathcal{N}(\mathbf{x}|0, \mathbf{1}) \quad \text{if } \theta \sim p(\theta). \quad (4.11)$$

In other words, a normalizing flow defines a parametric representation of a generic probability distribution $p(\theta)$, obtained by change of variables,

$$p_W^{\text{flow}}(\theta) = \mathcal{N}(\Phi_W(\theta)|0, \mathbf{1}) \mid \det J_{\Phi_W}(\theta) \mid \quad (4.12)$$

where J_{Φ_W} is the Jacobian of the flow transformation Φ_W . Sampling from p_W^{flow} can then be easily done by sampling $\mathbf{x} \sim \mathcal{N}(\mathbf{x}|0, \mathbf{1})$ and obtaining θ from the inverse flow transformation: $\theta = \Phi_W^{-1}(\mathbf{x})$. Thus, given a target distribution, both the problems of sampling and of density estimation become tractable thanks to the normalizing flow model.

The flow transformation Φ_W is built by *composing* n_{layers} simple (invertible) transformations, each called a layer. Of course, depending on the application, a variety of options are available in the literature. We build a layer by concatenating a linear transformation and a masked autoregressive layer [416–418] with n_{hidden} hidden features. A masked autoregressive layer implements the following transformation:

$$T_{\text{MADE}}(\theta) = a(\theta)\theta + b(\theta) \quad (4.13)$$

where the coefficients $a(\theta), b(\theta)$ are computed by (masked) autoencoders with n_{hidden} hidden features.

In our case, the target probability distribution has support in the rectangle $[\theta_{\min}, \theta_{\max}]$, while the base distribution of the flow (a Gaussian) has support in \mathbb{R}^D . We implement the change of support explicitly by introducing the following transformation $T_0(\theta) : [\theta_{\min}, \theta_{\max}] \rightarrow \mathbb{R}^D$ as the first layer of the flow:

$$T_0(\theta) = 0.5 \log \frac{1+y}{1-y} \quad \text{with} \quad y = \frac{2\theta - \theta_{\min} - \theta_{\max}}{\theta_{\max} - \theta_{\min}} \quad (4.14)$$

where the fraction above is intended as element-wise division.⁴ This transformation maps the rectangle $[\theta_{\min}, \theta_{\max}]$ into the plane. Then the remaining transformations only need to implement a change in probability density and not in the support of the distribution, making the loss function optimization easier.

The flow probability distribution $p_W^{\text{flow}}(\theta)$ is trained to closely reproduce a given probability distribution $p^{\text{target}}(\theta)$. During the training, the weights W of the flow are set by minimizing a loss function $\mathcal{L}_{\Phi}(W)$, which measures the discrepancy between p^{target} and p_W^{flow} . The minimization is performed by gradient descent. In our case, $p^{\text{target}} \propto \sqrt{\det M}$, with an unset normalization.

Depending on the nature of the data, several loss functions are available. If *samples* from the target distribution are available, the loss function is defined as the forward Kullback–Leibler (KL) divergence [419, 420] between the target distribution $p^{\text{target}}(\theta)$ and the one defined by the flow in Eq. (4.12):

$$\mathcal{L}_{\Phi}^{KL}(W) = -\mathbb{E}_{p^{\text{target}}(\theta)}[\log p_W^{\text{flow}}] + \text{const.} \quad (4.15)$$

where the expected value is computed using empirical samples from $p^{\text{target}}(\theta)$ to provide a Monte-Carlo estimation of the loss function.

In our situation however, we do not have access to such samples (indeed, we are training the flow precisely to avoid sampling!) but we are only able to evaluate p^{target} up to a constant scaling factor. For this reason, we treat the training as a *regression* problem, rather than a density estimation problem, and we use the following loss function:

$$\begin{aligned} \mathcal{L}_{\Phi}(W) &= \frac{1}{N} \sum_{i=1}^N \left(\log p_W^{\text{flow}}(\theta_i) - \log p^{\text{target}}(\theta_i) \right)^2 \\ &= \frac{1}{N} \sum_{i=1}^N \left(\log p_W^{\text{flow}}(\theta_i) - \log \sqrt{|M(\theta_i)|} + C \right)^2 \end{aligned} \quad (4.16)$$

where the sum runs on a dataset of N points:

$$\{(\theta_i, \sqrt{|M(\theta_i)|})\}_{i=1}^N. \quad (4.17)$$

Our experiments show that $N \simeq 5 \times 10^5$ is adequate in most cases.

In Eq. (4.16), C is a *trainable* constant, which sets the normalization of $p^{\text{target}} = e^{-C} \sqrt{|M|}$ on the domain of interest. Although not strictly needed, it can have a large impact on the flow performance, since it constrains the values of $\sqrt{|M(\theta)|}$ to a scale which is easier to learn by the normalizing flow. Some heuristics suggest initializing

⁴Note that the inverse T_0^{-1} of the transformation takes a simple form: $\frac{1}{2}[\tanh(T_0(\theta))(\theta_{\max} - \theta_{\min}) + \theta_{\max} + \theta_{\min}]$, where again the multiplication is intended as element-wise.

the constant to the 90th percentile of the values $\log \sqrt{|M(\theta)|}$ stored in the dataset. As shown in App. 4.C, the constant can be used to compute (an approximation to) the volume of the parameter space \mathcal{V} in Eq. (4.9).

The values of θ_i in the dataset Eq. (4.17) are obtained by sampling the masses m_1, m_2 from

$$p(\mathcal{M}_c, \eta) \propto \mathcal{M}_c^{10/3} \eta^{8/5}, \quad (4.18)$$

where $\mathcal{M}_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}$ is the chirp mass and $\eta = \frac{m_1 m_2}{(m_1 + m_2)^2}$ is the symmetric mass ratio. All other quantities are sampled from a uniform distribution in the coordinates.

Eq. (4.18) defines a flat distribution on the chirp-time parameters τ_0 and τ_3 [339]. Indeed, it can be shown that for a non-spinning binaries, the metric expressed in the chirp-time coordinates is approximately flat [325, 330], and that Eq. (4.18) represents a first order approximation to the true metric. Sampling from Eq. (4.18), ensures a high quality training set, where the distribution of the training points is reasonably close to the target distribution⁵.

During the training we halve the learning rate each time the validation loss does not improve more than a given threshold after a given number of iterations. This procedure finds local minima better in the loss function. We also apply early stopping, to avoid useless gradient descent iterations.

The training of the normalizing flow usually takes $\mathcal{O}(30$ minutes). On the other hand, from one to a few hours are needed to generate a dataset of $\mathcal{O}(10^5)$ points, depending on the dimensionality of the manifold and on the waveform approximant. This is the bulk of the cost of generating a template bank: the random template placing takes only a few minutes.

4.2.3 Random Template Placing

As customary, the input parameter controlling the average spacing and number of templates is the *minimal match* MM . It is defined as the minimum tolerable match that a random signal (inside the relevant parameter space) must have with its nearest templates in the bank. Of course, during the template placement, we only consider the match between templates on the same manifold, while the quantity can be used also to compare waveforms on different manifolds.

To generate our random template bank, following [345], we add random templates to the bank until a satisfactory coverage is achieved. The coverage is checked using a procedure that closely matches [346]. The templates are sampled from the

⁵Indeed more samples are present at low chirp mass, which is where the metric determinant tends to have larger values due to longer waveforms (for a constant starting frequency). Hence, a consistent bias in the low mass region is largely penalized in the loss function due to more samples in the dataset at low mass.

normalizing flow in Eq. (4.12), which, as discussed above, is trained to target Eq. (4.10). This choice makes sure that the templates are spread as “uniformly as possible” across the manifold.

One point of the space θ is said to be *covered* by the bank if there is at least one template θ_T in the bank, whose squared metric distance (mismatch) as given in Eq. (4.5) is at most $1 - MM$ or:

$$d_M^2(\theta, \theta_T) < 1 - MM. \quad (4.19)$$

The covering fraction $\hat{\eta}$ of a given region \mathcal{T} of the parameter space is then defined as the fraction of volume covered by the bank:

$$\hat{\eta}(\mathcal{T}) = \frac{1}{\text{Vol}(\mathcal{T})} \int_{\mathcal{T}} d^D \theta \sqrt{\det M(\theta_i)} c(\theta). \quad (4.20)$$

where $c(\theta)$ is an indicator function:

$$c(\theta) = \begin{cases} 1 & \text{if } \theta \text{ is covered by the bank} \\ 0 & \text{otherwise} \end{cases}. \quad (4.21)$$

According to our chosen placement method, we do not require that the space is fully covered but we only require that it is covered with probability η . This means that we terminate the bank construction when the covering fraction $\hat{\eta} \geq \eta$.

To provide a sensible estimate of the covering fraction $\hat{\eta}$, we perform a Monte Carlo estimation of the integral in Eq. (4.20) [346]:

$$\hat{\eta}(\mathcal{T}) \simeq \frac{1}{N_{\text{livepoints}}} \sum_i c(\theta_i) \quad (4.22)$$

where the $N_{\text{livepoints}}$ samples $\theta_i \sim p^{\text{flow}}$ are sampled from the normalizing flow and are called *live points*. Note that in Eq. (4.22), we don’t compute volumes using the volume element $\sqrt{\det M}$ itself but rather its normalizing flow approximation.

In practice, while the templates are being added to the bank, the distance between each livepoint is computed. If the i -th livepoint is close enough to the newly added template, it will be removed from the set of live points and a running estimate of $\hat{\eta}(\mathcal{T})$ will be updated. The estimation of the covering fraction $\hat{\eta}$ has standard deviation [346, App. A]:

$$\sigma_{\hat{\eta}} = \sqrt{\frac{\eta(1 - \eta)}{N_{\text{livepoints}} - 1}}, \quad (4.23)$$

which suggests using a large number of live points for better estimation. In [346], the authors typically choose $\eta = 0.9$ and $N_{\text{livepoints}} = 2000$.

Since the method does not check for distances between templates, it can overcover the space (as also reported in [345, 346]), especially for a low number of dimensions.

Despite this, it is very fast and provides a reliable bank at a cheap computational and memory cost. Moreover, as argued in [344, 345, 347], for a large number of dimensions, the banks generated by the random method provides close to optimal performance.

As a final remark, we note that for the purpose of computing the covering fraction, the templates do not need to be stored, which enables the algorithm to run with a very low memory footprint. As exemplified in Sec. 4.5.3, this allows to study the number of templates required to cover a particular region of the parameter space, providing invaluable pieces of information useful to plan a GW search.

4.3 Validation

In this section, we assess the performance of the two key ingredients of our template bank generation algorithm, namely the normalizing flow model and the random placement algorithm. Our goal is to understand the limitations of our algorithm as well as to make an informed choice of the various hyperparameters that impact the quality of the template bank.

We will consider different manifolds, which will be named with a string that lists the manifold coordinates. The coordinates are grouped by mass coordinates, spin coordinates, (eventual) eccentricity coordinates (i.e. e and a) and (eventual) angles coordinates (i.e. ι and ϕ). Consequently, a string has the format:

`Masses_Spin1_Spin2_Eccentricity_Angles.`

Valid options for the mass coordinates are `m1m2` which uses m_1 and m_2 as coordinates, `Mq` which uses total mass $M = m_1 + m_2$ and mass ratio $q = m_1/m_2 > 1$, and `logMq` which uses $\log_{10} M$ instead of M . Similarly, other variables are listed by their names. The manifold with spin label `chi` uses the effective spin χ_{eff} parameter as coordinate. Since χ_{eff} is degenerate in the two spins, we choose to set $s_{1z} = s_{2z} = \chi_{\text{eff}}$ and all the other spin components to 0.

If more than one spin coordinate is given for a given BH, the spin vector \mathbf{s} will be parameterized in spherical coordinates with magnitude $s \in [0, 1)$ and angles $\theta \in [-\pi, \pi]$ and $\varphi \in [0, \pi]$ as follows:

$$s_x = s \sin \theta \cos \varphi \quad (4.24)$$

$$s_y = s \sin \theta \sin \varphi \quad (4.25)$$

$$s_z = s \cos \theta. \quad (4.26)$$

Note that the angle θ controls the amount of precession. With $\theta = 0, \pm\pi$ the spin has only a z component (i.e., is aligned with the orbital angular momentum), while

	Parameter space	D	Architecture
m1m2_nonspinning	$m_1, m_2 \in [1, 200] M_{\odot}$ $q \in [1, 30]$ $f \in [15, 1024] \text{ Hz}$ IMRPhenomD [157]	2	60 60 30
Mq_s1xz	$M \in [25, 100] M_{\odot}$ $q \in [1, 5]$ $s_1 \in [0, 0.99]$ $\theta_1 \in [0, \pi]$ $f \in [15, 1024] \text{ Hz}$ IMRPhenomXP [47]	4	70 70
m1m2_chi_e	$m_1, m_2 \in [1, 50] M_{\odot}$ $q \in [1, 20]$ $\chi_{\text{eff}} \in [-0.99, 0.99]$ $e \in [0, 0.5]$ $f \in [10, 1024] \text{ Hz}$ EccentricFD [421]	4	60 60 60
logMq_s1z_s2z_iota (with HM)	$m_1, m_2 \in [50, 300] M_{\odot}$ $M \in [100, 400] M_{\odot}$ $q \in [1, 10]$ $s_{1z}, s_{2z} \in [-0.99, 0.99]$ $\iota \in [0, \pi]$ $f \in [10, 1024] \text{ Hz}$ IMRPhenomXP [47]	5	20 60 60
logMq_s1xyz_s2z_iota	$m_1, m_2 \in [1, 100] M_{\odot}$ $M \in [2, 150] M_{\odot}$ $q \in [1, 20]$ $s_1 \in [0, 0.99]$ $\theta_1 \in [-\pi, \pi]$ $\varphi_1 \in [0, \pi]$ $s_{2z} \in [-0.99, 0.99]$ $\iota \in [0, \pi]$ $f \in [15, 1024] \text{ Hz}$ IMRPhenomXHM [422]	7	100 60 60 60

Table 4.1: Details of the manifold considered for the validation of the normalizing flow model in Fig. 4.1. For each manifold, we report the variables being sampled together with their ranges. We also list the frequency range considered, the waveform approximant used, the number of dimensions D of the manifold as well as the number of hidden features for each layer of the flow.

for $\theta = \pm\pi/2$ there is maximal precession, as the spin vector only has an in-plane component.

4.3.1 Normalizing Flow Validation

To study the accuracy of the normalizing flow model in reproducing $\sqrt{|M|}$, we consider five manifolds. The manifolds are listed in Tab. 4.1, together with the region of the parameter space they cover. We also report the waveform approximant used as well as the frequency range where the metric is computed. The manifolds were chosen to have a variety of number of dimensions D and to cover a broad ranges of physical scenarios (non-spinning, aligned-spins, precession, HM, and eccentric orbits).

For each manifold we generate a dataset of 3×10^5 points and we compute the (log) value of the PDF in Eq. (4.10). We then train a normalizing flow model on each of the datasets. The architecture of each flow is also reported in Tab. 4.1. Fig. 4.1 shows a histogram with the accuracy of the normalizing flow reconstruction of the PDF on each manifold. This is quantified by $\log_{10} \frac{p^{\text{flow}}}{p^{\text{true}}}$, which measures the logarithmic ratio between the two PDFs.

Overall, the accuracy of the flow is (almost) always contained within one order of magnitude. Whether a similar error is acceptable for the purpose of template placement needs to be checked on a case-by-case basis with an injection study, as discussed in Sec. 4.4.

We note that all histograms are well-centered around 0, showing that the flow does not have a systematic bias. Moreover, the accuracy tends to be higher for low-dimensional manifolds. Indeed, low dimensional manifolds present an easier learning task for the flow.

The manifold `logMq_s1xyz_s2z_iota` shows the largest spread in accuracy, as it is the largest dimensional manifold being considered. Note that it parameterizes a huge parameter space, which cannot be realistically covered by a template bank. Hence, as a realistic bank will necessarily cover a subset of the manifold, a flow trained on that smaller parameter space will most certainly show better accuracy, due to an easier regression task.

Finally, we see that the flow trained on the eccentric manifold `m1m2_chi_e` has remarkably good performance. This can be explained by the fact that the approximant `EccentricFD` [421] used is analytical. This ensures very smooth behaviour across the parameter space, which can be easier for the normalizing flow model to learn.

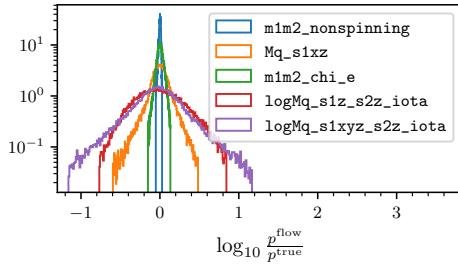


Figure 4.1: Study of the accuracy for several normalizing flow, trained on different manifolds. For each manifold, we compute the logarithmic ratio $\log_{10} \frac{p^{\text{flow}}}{p^{\text{true}}}$ between the PDF computed by the flow and the true one. We use 40000 test points from the validation set of each manifold. Details on the manifold considered are reported in Tab. 4.1.

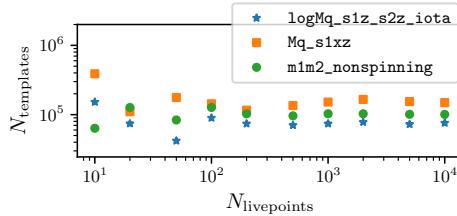


Figure 4.2: Validation of the random template placement algorithm. For three of the manifolds introduced in Tab. 4.1, we plot the number the number of templates $N_{\text{templates}}$ of a random template bank as a function of the number of live points $N_{\text{livepoints}}$ used to estimate the covering fraction. For each template bank, we set $\eta = 0.9$ and $MM = 0.97$.

4.3.2 Template Placement Performance

As already stated, the template placement method in use closely matches the one introduced in [346]. The main novelty introduced here is sampling with the normalizing flow as opposed to rejection sampling.

For the random placement method, there are two parameters to tune that affect the final bank size. They are the number of live points $N_{\text{livepoints}}$ and the covering fraction η . The authors of [346] make an extensive investigation on how the bank size depends on such quantities and we do not repeat such in-depth studies here. We limit ourselves to examining the convergence of the template number $N_{\text{templates}}$ as a function of $N_{\text{livepoints}}$ (see [346, Fig. 4 (right)]) in the case of manifolds with precessing and HM signals. For the study, we chose the manifolds `m1m2_nonspinning`, `Mq_s1xz` and `logMq_s1z_s2z_iota` introduced in Sec. 4.3.1 (see also Tab. 4.1). The second manifold covers a precessing parameter space, while the metric on the latter

	Parameter space	Size	
		Original	<code>mbank</code>
HM bank [328]	$M \in [50, 400] M_{\odot}$		
	$q \in [1, 10]$		
	$\iota \in [0, \pi]$	20500	58932
	$\phi \in [0, 2\pi]$		
IMRPhenomXHM [422]			
"All-sky" bank [380]	$m_1, m_2 \in [1, 200] M_{\odot}$		
	$q \in [1, 20]$		
	$\chi_{\text{eff}} \in [-0.99, 0.99]$	1.8×10^6	1.3×10^6
	IMRPhenomD [157]		

Table 4.2: Details of the two banks available in the literature that we reproduce with our code. For each bank, we indicate the parameter space considered and the approximant used. We also compare the number of templates of the banks obtained with the different methods.

manifold is computed with an HM approximant [422].

We present our results in Fig. 4.2, where the number of templates is computed with a covering fraction $\eta = 0.9$ with varying $N_{\text{livepoints}}$. In all cases the number of templates converges to a constant value as $N_{\text{livepoints}}$ increases. Already ~ 500 live points are enough to provide an accurate estimation of the bank size. Our results are consistent with the findings of [346], which we further extend to higher-dimensional manifolds.

4.4 Comparison with Other Bank Generation Methods

We compare the output of `mbank` with two banks available in the literature, generated with two different methods. The first bank is a non-spinning HM bank [328], covering the high mass region of the BBH parameter space. The bank was generated using the stochastic placement algorithm, as implemented in the code `sbank` [351]. The second bank is the aligned-spin bank [380] currently in use by the `GstLAL` pipeline [77,78] for the fourth observing run (O4) of the LIGO-Virgo-KAGRA collaboration. It was generated using the `manifold` [343] metric template placement algorithm called and covers a very wide mass range in the BNS and BBH parameter space. Both banks have a minimal match MM requirement of 0.97.

In much of what follows we will measure the coverage of a bank by studying the dependence of the fitting factor Eq. (3.45) across different regions of the parameter space. To do so, we randomly extract a number of simulated signals and, for each of them, we compute the maximum match with the templates of the bank.

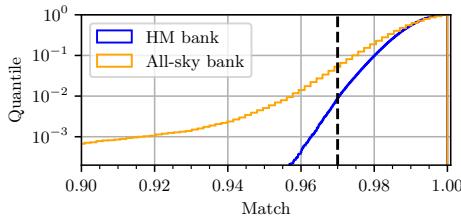
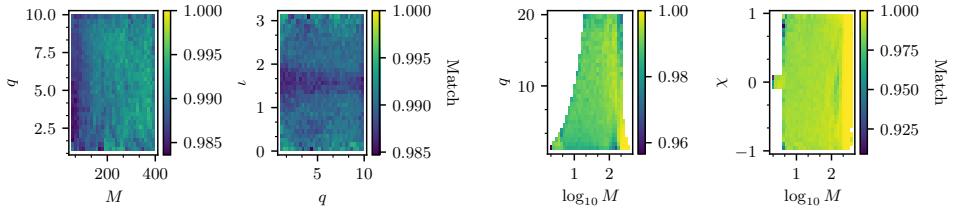


Figure 4.3: Fitting factor studies for the two template banks introduced in Sec. 4.4. As discussed in Sec. 4.4.1 and 4.4.2 respectively, “HM bank” is designed to reproduce [328] and targets high mass non-spinning systems with HM content, while the “All-sky bank” bank covers aligned-spin systems (without HM) over a broad mass range, following [380]. We report the cumulative histogram of the fitting factors of 10^5 injections samples across the parameter space.



(a) Validation of the “HM bank”, generated with our code and designed to reproduced [328]. (b) Validation of the “All-sky bank”, generated with our code and designed to reproduced [380].

Figure 4.4: Validation of the two template banks designed to reproduce two independently generated banks. On the left side, we report results pertaining the “HM bank”, while on the right side we report the “All-sky bank”. For each two dimensional bin, we report the median fitting factor of 10^5 injections covering the parameter space, as described in the text.

4.4.1 A Non-spinning HM Template Bank

The non-spinning HM bank described in [328] covers systems with total mass M in the range $[50, 400]M_\odot$ and mass ratio $q \in [1, 10]$. It also includes the inclination angle ι and reference phase ϕ of the system, both covering the whole possible spectrum of values $\iota \in [0, \pi]$ and $\phi \in [0, 2\pi]$. The authors use the analytical “zero-detuning high power” PSD [423] and consider a low frequency cutoff $f_{\min} = 10$ Hz.

As already noted, they use the state-of-the-art code `sbank` [349,351]. The method is very accurate and known to provide effective coverage with a low number of templates. Of course, this comes at a large up-front computational cost to construct the bank.

To reproduce this bank, we place templates on the manifold `logMq_nonspinning--iotaPhi`, with coordinates $\log_{10} M$, q , ι and ϕ . We use the same PSD and coordinate ranges as the original bank. We refer to our bank as “HM bank”. We train a normalizing flow model with 4 layers with 60, 60, 60, 10 hidden features respectively and we choose $N_{\text{livepoints}} = 2000$ and a covering fraction $\eta = 0.8$. Our bank has 58932 templates and took a few hours to generate; the original bank is reported to have 20500 templates. All information is summarized in Tab. 4.2. We perform an injection study, drawing 10^5 signals uniformly sampled in $\log M$, q , $\cos \iota$ and ϕ . The results of such study are reported in Fig 4.3 and Fig 4.4a.

First we note that our bank successfully covers the parameter space, with only 1% of injections found with fitting factor below 0.97 and less than 1% with fitting factor below 0.96. The coverage of the bank is similar to that of [328]. In Fig. 4.4a, we observe that the coverage is uniform across the space, i.e. we do not see regions where the fitting factor is significantly different from the others.

Comparing the number of templates, it is striking that our bank has almost three times more templates than the original template bank. As no template rejection is done during the random bank construction, there is no control over templates being too close to each other. For this reason, an over-coverage of the space is inherent to the random template placement and is also reported in [345, 346]. This problem can be addressed in future work, as discussed in Sec. 4.6.

4.4.2 An “All-sky” Template Bank

The aligned-spin bank (with no HMs) introduced in [380] covers a broad mass range, with systems with component masses $m_1, m_2 \in [1, 200] M_\odot$. The spins of the two objects are constrained to be equal to each other⁶, $s_{1z} = s_{2z} = \chi_{\text{eff}}$, spanning the range $[-0.99, 0.99]$. The authors set an upper limit to the mass ratio $q < 20$. Moreover, for objects with component mass $m < 3 M_\odot$, they limit χ_{eff} in the range $[-0.05, 0.05]$ ⁷. The authors use the Advanced LIGO O4 Design PSD (with 190 Mpc range) [425] and consider a low frequency cutoff $f_{\text{min}} = 10 \text{ Hz}$.

The comparison with [380] is particularly interesting, since the bank is also produced with a metric template placement, implemented in the `manifold` code [343]. `manifold` uses a geometric approach, where the parameter space is iteratively split into (hyper)rectangles along the coordinates, until the volume of each rectangle reaches a sufficiently small value that it can be covered by a single template.

As summarized in Tab. 4.2, we construct a bank to cover the parameter space

⁶This choice reduces the dimension of the manifold, without compromising the template bank accuracy.

⁷This is motivated by astrophysical considerations. Objects with masses smaller than $3 M_\odot$ are likely to be neutron stars and such objects are believed to develop only mild rotations [424].

used in [380] over the manifold `m1m2_chi`, sampling the coordinates m_1 , m_2 and χ_{eff} . We may refer to our bank as the “All-sky” bank. To produce our “All-sky” bank, we trained three different normalizing flows in different regions of the parameter space. A first normalizing flow covers the BBH region with $m_1 \in [3, 200] \text{ M}_\odot$, with $\chi_{\text{eff}} \in [-0.99, 0.99]$. A second one covers the BNS region, covering the manifold, $(m_1, m_2, \chi_{\text{eff}}) \in [1, 3] \text{ M}_\odot \times [1, 3] \text{ M}_\odot \times [-0.05, 0.05]$. A third normalizing flow specializes in the high mass region, characterized by $m_1, m_2 \in [100, 200] \text{ M}_\odot$. Indeed, at high masses, the template density is so low that hardly any livepoint is sampled, which results in dramatic undercoverage. An appropriate coverage is enforced by the third normalizing flow, which places $\mathcal{O}(3000)$ templates in the region as opposed to *zero* templates placed by the first flow. The additional coverage at high masses is manifest in Fig. 4.4b, as discontinuity in the fitting factor for $m_1, m_2 > 100 \text{ M}_\odot$.

All the three normalizing flow models are made of 5 layers of 10 hidden features each. Three templates banks are generated using each normalizing flow and they are merged together afterwards. For the template placement we set $N_{\text{livepoints}} = 2000$ and covering fraction $\eta = 0.95$. The resulting bank has 1326805 templates.

The bank generation took around three hours, with most of the computing time spent on the dataset generation (i.e. on expensive metric evaluations). If needed, the dataset generation can be easily parallelized using `mbank`, hence reducing significantly the bank generation time. Relying on parallel execution, [380] reported a generation time of minutes.

To validate our bank, we generate an injection set with 10^5 injections, with the logarithm of the masses uniformly sampled. Results of our injections studies are reported in Fig 4.3 and Fig 4.4b. Note that our injection set is different from the ones used in [380].

In Fig 4.3, we see that $\sim 5\%$ of the injections have a match below 0.97. The low fitting factor injections are mostly located around the low mass corners of the bank, clustered on the low mass end of the BNS region and in the high spin - low mass edge of the BBH region. Inside the template bank and on the high mass end of the parameter space, satisfactory coverage is achieved. Our results suggest that `mbank` struggles to accurately cover the “narrow” corners of the parameter space. Nevertheless, this is a common problem that has been observed with other placement methods as well, and several strategies have been proposed to cope with it. Within our framework, the simplest option would be to extend the boundaries of the bank at low masses, thus ensuring better coverage of the region of interest.

With slight variations depending on the region of parameter space, [380] reports that 10% of BBH injections have fitting factor smaller than ~ 0.98 , while for our bank the 10th percentile is around 0.975. Even though it is hard to compare the results

directly due to different injection sets, it seems fair to state that, compared to [380], our template bank provides slightly worse injection recovery. On the other hand, our template bank has 30% *less* templates, matching the number of templates placed by `sbank` in the same region, as reported by [380]. With an accurate treatment of the low mass corner, the coverage of our template bank will easily match the one of [380], with a comparable bank’s size.

4.5 Novel Applications of the Method

Our template placement method allows for several exciting applications in GW data analysis. Obviously, the most straightforward application is the generation of high-dimensional template banks, such as a precessing and/or HM banks. While in principle it is possible to generate these high-dimensional banks with a stochastic placement method, very few of such banks have been generated so far, mostly due to the enormous computational cost of choosing the right parameter space and of computing the match between templates. These tasks becomes feasible thanks to `mbank`.

Besides efficient high-dimensional bank generation, our method can be used for other purposes as well. These include choosing the appropriate parameter space to cover by forecasting the size of a bank or selecting the appropriate coordinates to cover a given region of binary systems. Moreover, our normalizing flow could be used as a proposal for a stochastic placement algorithm or to create datasets for machine-learning applications in GW data analysis.

In what follows, we generate a large precessing template bank and a large aligned-spin HM bank. Additionally, we provide a detailed discussion of other innovative applications of our method.

4.5.1 A Precessing Bank

Choosing the Parameter Space

The main difficulty in generating a precessing bank lies in the huge size of the parameter space. As we show below, a precessing bank can easily have *billions* of templates, even when covering the mass range routinely explored by “standard” searches. As current search pipelines can handle only up to a few *million* templates, due to computational cost limitations, the size of a bank sets very stringent constraints in the selection of a suitable parameter space to explore with a GW search.

Another difficulty, related to the first, arises from the choice of the BBH coordinates to include in the bank, i.e. the choice of manifold. In principle, a precessing BBH system is described by 10 parameters (two masses, six spins, and two angles).

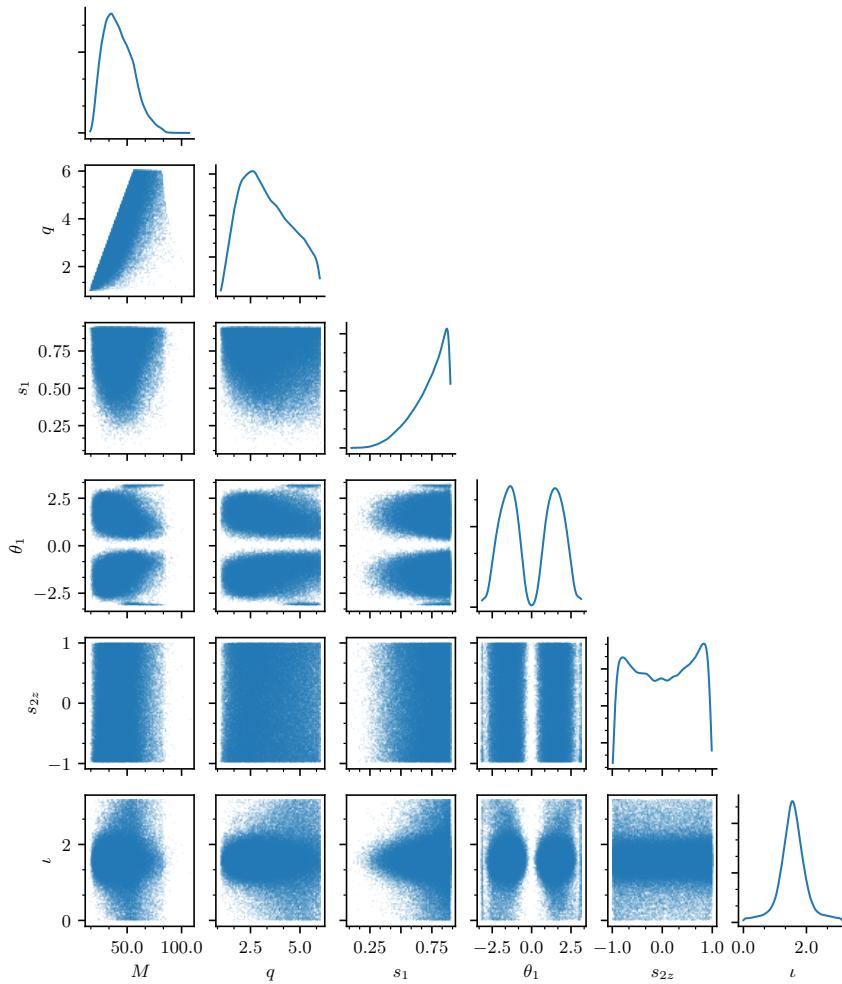


Figure 4.5: Corner plot with the templates of the precessing bank described in Sec. 4.5.1. Along the diagonals, we show the histogram of the template number as a function of each coordinate.

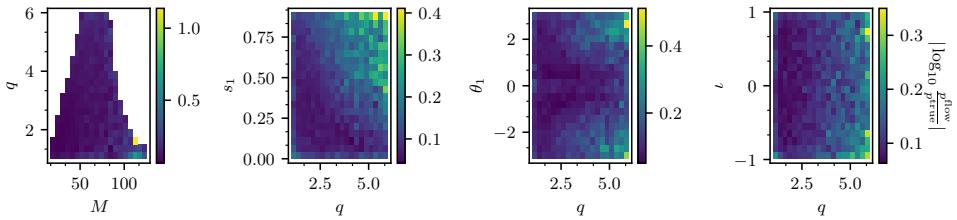


Figure 4.6: Accuracy of the normalizing flow trained used to generate the precessing bank in Sec. 4.5.1. The accuracy is expressed in terms of the logarithmic ratio between the template density PDF p^{true} Eq. (4.10) and its approximation p^{flow} given by the flow. The flow accuracy is evaluated on 40000 test points.

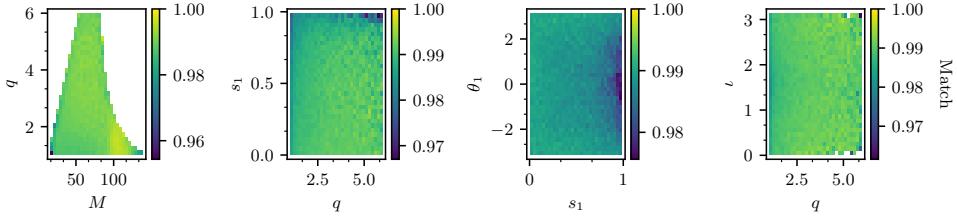


Figure 4.7: Fitting factor study of the precessing bank, introduced in Sec. 4.5.1. For each bin, we color-code the median fitting factor of 10^5 injections sampled “On manifold”, as described in the text.

However, not all of them are important, as large changes in some parameters do not result in large changes in the waveform morphology. Thus, including them in the bank does not yield any obvious improvement and, on the contrary, it may lead to vanishing metric eigenvalues, which would degrade the metric predictivity, hence the template placement. The latter point is discussed with more details in Sec. 4.5.3.

Finally, a more technical complexity arises from the fact that in high dimensional spaces, both the training of a normalizing flow (see Sec. 4.3.1) and the template placement become harder, hence possibly harming the quality of the template bank.

All these difficulties imply that great care must be taken when deciding both the parameter space and the BBH variables to include in the bank. The choices are entangled, since covering different manifolds with the same mass range can produce banks of very different sizes. Roughly speaking, choosing a lower dimensional submanifold reduces the bank size, at the cost of a loss in the bank’s ability to cover the high dimensional space.

To choose a manifold, we rely on the theory. In [183], the authors find that the

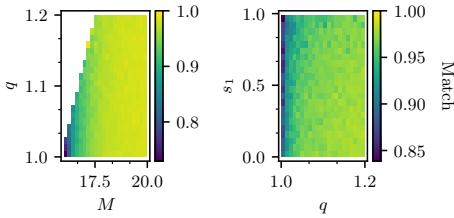


Figure 4.8: Fitting factor study of the precessing bank of Sec. 4.5.1. Unlike Fig. 4.7, here we focus on the low q , low M region, where the random placement method fails. For each bin, we color-code the median fitting factor of 5×10^4 injections sampled “On manifold”, as described in the text.

effect of the four in-plane spin components (i.e. $s_{1x}, s_{1y}, s_{2x}, s_{2y}$) can be well approximated by a single precessing spin parameter χ_P assigned to the x -component of the heavier object’s spin. Thus, a generic precessing system is roughly equivalent to a system with

$$\begin{aligned} \mathbf{s}_1 &= (\chi_P, 0, s_{1z}) \\ \mathbf{s}_2 &= (0, 0, s_{2z}) \end{aligned}$$

effectively creating an explicit mapping between a six dimensional spin manifold to a three dimensional one. In a later work [186], it is suggested that to capture the combined effect of precession and HM, a two-dimensional spin parameter $\vec{\chi}_P$ is needed. In this case, the mapping is between a six-dimensional spin manifold to a four-dimensional one.

Both works suggest that the in-plane components of the spin on the lighter object (i.e. s_{2x}, s_{2y}) can be neglected, reducing the dimensionality of the parameter space. Moreover, since we are not currently concerned with precession combined with HM⁸, we can rely on the one-dimensional effective spin mapping [183] to also neglect the y -component of the spin of the heavier object, s_{1y} .

We then consider only three out of six spin components, s_{1x}, s_{1z} and s_{2z} , where all the effects of precession are included in s_{1x} . To obtain accurate coverage, we also need to include the inclination ι in the manifold. Some investigations showed that the inclusion of the reference phase ϕ yields a (almost) degenerate metric, which, by dramatically undercovering the space, negatively affects the placement. Luckily, as injection studies show that neglecting ϕ does not harm the bank’s effectiveness, we can exclude ϕ from the set of parameters. However, this might not be the case if we include both precession and HMs.

⁸In such a space, the template banks would be unfeasibly large!

To summarize, we find that the 6 variables M , q , s_{1x} , s_{1z} , s_{2z} and ι provide a sufficiently complete description of waveforms in the precessing space. This claim is confirmed by an injection study presented in Fig. 4.9, where we see that more than 93% of the injections covering the 10 dimensional precessing space have a fitting factor greater than the minimal match target of 0.97. We note that a precessing template bank with HMs will likely need to sample two additional variables s_{1y} and ϕ , hence increasing the dimensionality to 8 [186].

Regarding the search parameter space, we are interested to target BBHs where precession is stronger as such systems are most likely to be missed by current searches [74, 185]. Precession is more visible for high mass ratio, edge-on⁹ systems and for high values of spins [73]. Moreover, as more cycles are detectable, precession effects will be stronger for longer signals due to the accumulation of the phasing effects of precession. These considerations suggest that very asymmetric, low mass systems, such as the neutron star-black hole (NSBH) space, would be an ideal target for a precessing bank. However, as shown below in Sec. 4.5.3 searching the full NSBH region is unfeasible, as hundreds of millions of templates would be needed.

For this reason, we restrict ourselves to a different, less extreme, region of the parameter space. After several investigations, made possible by the speed and flexibility of our approach, we found that a parameter space with component masses in the range $[8, 70] M_{\odot}$, with a mass ratio cut-off of 6, produces a bank with a manageable size. In this space, we obtain a precessing bank with ~ 2 million templates. Extending the parameter space to lower masses (or higher mass ratios) results in much larger banks, pushing the limits of current pipelines.

In closing, we stress again that the investigations above are made possible by `mbank`, since they rely on fast template bank generation across a variety of manifolds and ranges of coordinates.

Generating and Validating the Bank

As stated above, our precessing bank covers the manifold `logMq_s1xz_s2z_iota`, with coordinates $\log_{10} M$, q , s_1 , θ_1 , s_{2z} and ι . We consider BBHs with individual masses between 8 and $70 M_{\odot}$, with a maximum mass ratio $q = 6$. The other variables s_1 , θ_1 , s_{2z} and ι cover the set $[0, 0.9] \times [-\pi, \pi] \times [-0.99, 0.99] \times [0, \pi]$.

To compute the metric, we use the Advanced LIGO O4 sensitivity estimate [425] and we set a frequency range of $[15, 1024]$ Hz, employing the approximant `IMRPhenomXP` [47]. We train a normalizing flow with 3 layers with 100, 100 and 60 hidden features respectively, using a dataset of 4×10^5 points. The flow performance after training is reported in Fig. 4.6. To generate the bank, we use a minimal match requirement of

⁹An edge-on system is observed with inclination $\iota \simeq \pi/2$.

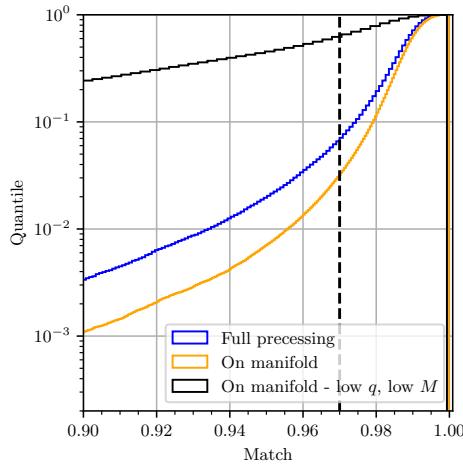


Figure 4.9: Cumulative fitting factor for the precessing bank introduced in Sec. 4.5.1. The 10^5 injections “Full precessing” have isotropic spins, while the 3×10^5 precessing injections “On manifold” are sampled on the manifold `logMq_s1xz_s2z_iota` and they have $s_{1y} = s_{2x} = s_{2y} = \phi = 0$. For the injections “On manifold”, we plot separately the low q , low M corner, characterized by $q \leq 1.2$ and $M \leq 20 M_\odot$. The other two histograms exclude this region.

0.97, with a covering fraction $\eta = 0.95$, estimated with 3000 live points. In a similar way to what was done for the “All-sky” template bank, we also train a normalizing flow to target the high total mass region with $M > 100 M_\odot$. We use the latter to place templates with the same covering fraction $\eta = 0.95$, with great benefits. The overall bank has 1605625 templates, plotted in Fig. 4.5.

This bank generation took a few hours in total: ~ 1 hour for the dataset generation, ~ 30 minutes for the training of the flow and ~ 5 minutes for the template placing. All the steps above ran on a single core, using less than 4 GB of memory. We highlight that our time and memory requirements are a fraction of those of a similar bank with the state-of-the-art stochastic algorithm.

The template distribution reported in Fig. 4.5 shows a spike in the template density for $\theta_1 = \pm\pi$ (close to the non-precessing limit) in the high mass ratio and high s_1 region. Some investigations indicate that these are not artifacts introduced by the normalizing flow. Whether the feature is physical or is due to the behaviour of the waveform approximant in the non-precessing limit remains an open question which needs more inspection.

To study the performance of our template bank, we generate two injections sets, with masses sampled uniformly in $\log m_1$ and $\log m_2$. The first set, labeled “Full precessing” has fully precessing injections (with two 3D spins and varying ϕ). The sec-

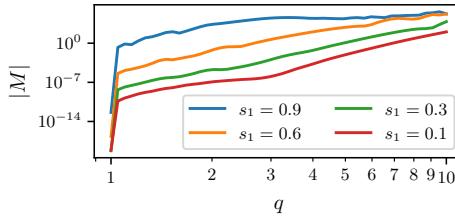


Figure 4.10: Determinant of the metric $|M|$ as a function of mass ratio q for different values of s_1 . The metric is evaluated on the manifold `Mq_s1xz_s2z_iota`, with $M = 10M_\odot$, $\theta_1 = \pi/2$, $s_{2z} = -0.3$ and $\iota = \pi/2$. It is manifest that in all cases, the metric determinant vanishes while $q \rightarrow 1$.

ond one, denoted as “On manifold”, has injections lying on the manifold `logMq_s1xz_s2z_iota`, hence covering a subset of the “Full precessing” set. The latter set is needed to asses the coverage of the bank on the manifold on which the templates lie and thus is a measure of the templates’ placement accuracy. On the other hand, the “Full precessing” injection set evaluates the ability of the bank to recover a generic precessing signal, hence assessing the quality of our choice of manifold. Clearly, this is the injection set that is most relevant for designing the bank for a fully precessing search.

We report the results of our study in Fig 4.9, in the form of a histogram of the fitting factors, and in Fig. 4.7, where we study the dependency of the fitting factor across the parameter space. Fig. 4.8 reports the same fitting factor study focused on the low q , low M region.

As is clear from Fig. 4.8 and 4.9, the random template placement method *fails* for the low q , low M region, with $q \leq 1.2$ and $M \leq 20M_\odot$, where only $\sim 40\%$ of the injections “On manifold” have a fitting factor higher than 0.97. On the other hand, outside the low q , low M corner, the template bank provides a good coverage: 97% of the injections “On manifold” has a fitting factor large than 0.97.

The poor performance for low mass ratio and low masses was also observed in the “All-sky” template bank in Sec. 4.4.2, although less severe. Such failure be explained by two combined causes. First of all, as noted above, the random method is unable to cover “sharp” corners of the parameter space, due to the lack of appropriate boundary treatment: this can (and does) severely limit the bank’s ability to cover the space. Moreover, we observe that for $q \rightarrow 1$ the metric determinant goes rapidly to 0, meaning that very few templates are placed. This is shown in Fig. 4.10, where we plot $|M|$ as a function of q keeping constant all the other coordinates¹⁰. The two effects com-

¹⁰Although not reported here, the same behaviour is observed for “standard” signals.

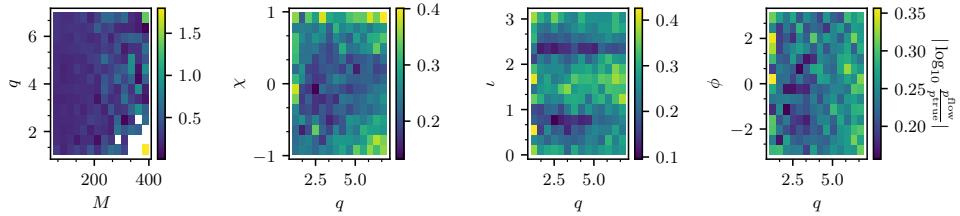


Figure 4.11: Accuracy of the normalizing flow trained used to generate the aligned-spin HM bank of Sec. 4.5.2. The accuracy is expressed in terms of the logarithmic ratio between the template density PDF p^{true} Eq. (4.10) and its approximation p^{flow} given by the flow. The flow accuracy is evaluated on 40000 test points.

bines together in the low q , low M region, which is drastically undercovered. The same issue is not observed anywhere else in the parameter space.

In principle, we could remedy the problem by extending the covered region to lower masses and higher q : this would make sure that the low q , low M target region does not lie at the boundaries of the bank anymore. However, the lack of coverage in this region is not a major concern for the bank’s effectiveness in a real search scenario. Indeed, precession for $q \sim 1$ has very little effect on the BBH waveform and a precessing system with symmetric masses would likely be detected by current aligned-spin searches.

In Fig. 4.7, we see that the coverage is rather uniform across the parameter space. The median fitting factor slightly drops for the high q high s_1 corner of the parameter space. As shown in Fig. 4.6, the flow performance degrades in that undercovered corner of the space: the true template density $\sqrt{|M|}$ is underestimated by the normalizing flow, which accordingly places less templates than optimal.

The fitting factor of the “Full precessing” injection set is fairly good, with only 7% of the injections (outside the “low q , low M ” region) below the target match. This means that the χ_P approximation that motivates our choice is robust: the manifold `logMq_s1xz_s2z_iota` provides a faithful low-dimensionality representation of the entire precessing parameter space.

4.5.2 An Aligned-spin HM Bank

Aligned-spin HM template banks are easier to generate than precessing ones, due to a smaller dimensionality of the parameter space. Indeed, a generic aligned-spin binary system with HMs is characterized by 6 parameters (two masses, two spins and two angles ι, ϕ) but, as for the non-HM case, the spin effects can be easily parameterized with an effective spin parameter, reducing the number of dimensions to 5. Note that

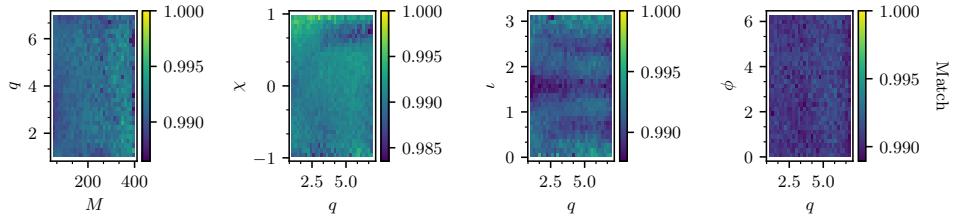


Figure 4.12: Fitting factor study of the aligned-spin HM bank, introduced in Sec. 4.5.2. For each bin, we color-code the median fitting factor of 10^5 injections sampled uniformly from the parameter space.

here we deal with one dimension more than in the non-spinning HM bank produced in Sec. 4.4.1. Despite less uncertainties in the choice of manifold than in the precessing case, the parameter space is very large and producing a template bank of a feasible size still requires a careful choice of the region to target.

We used `mbank` to generate an HM aligned-spin bank, covering the high mass region of the BBH parameter space. High mass events are notoriously hard to detect [371, 372]. As they are very short, their morphology matches closely non-Gaussian transient noise bursts, also called *glitches*, [211–214]. In this scenario, a more realistic model for the waveform can improve the detectability of such signals, thanks to both an increase in recovered SNR and to a more accurate signal-based veto [77, 426]. Several studies [64, 67, 407, 427] confirmed this claim, finding that failing to consider HMs in GW searches can lead to a large sensitivity loss for large mass ratios $q \gtrsim 4$ and high masses $M \gtrsim 100 M_\odot$ [73].

Consequently, our bank covers the manifold `logMq_chi_iotaphi`, sampling $\log_{10} M$, q and χ_{eff} as well as inclination and reference phase. We consider templates with total mass M between $50 M_\odot$ and $400 M_\odot$ and a mass ratio smaller than 7. The effective spin lies in range $[-0.99, 0.99]$ and, as usual, $\iota \in [0, \pi]$ and $\phi \in [-\pi, \pi]$. We use the Advanced LIGO O4 sensitivity estimate [425] and we set a frequency range of $[10, 1024]$ Hz, with approximant `IMRPhenomXHM` [422].

We generate a dataset with 4×10^5 points and train a normalizing flow with 4 layers, each with $n_{\text{hidden}} = 60$ hidden features. The accuracy of the normalizing flow is reported in Fig. 4.11. For the template placement, we use a minimal match requirement of 0.97 and set a covering fraction $\eta = 0.8$, estimated with 10000 live points. The overall bank gathers 2115299 templates, which are plotted in Fig. 4.13. The bank generation took roughly the same time as for the precessing bank.

We study the bank performance with 10^5 injections and report their fitting factor in Fig. 4.12 and Fig. 4.14. Our injection study shows that only $\sim 2\%$ of the injections

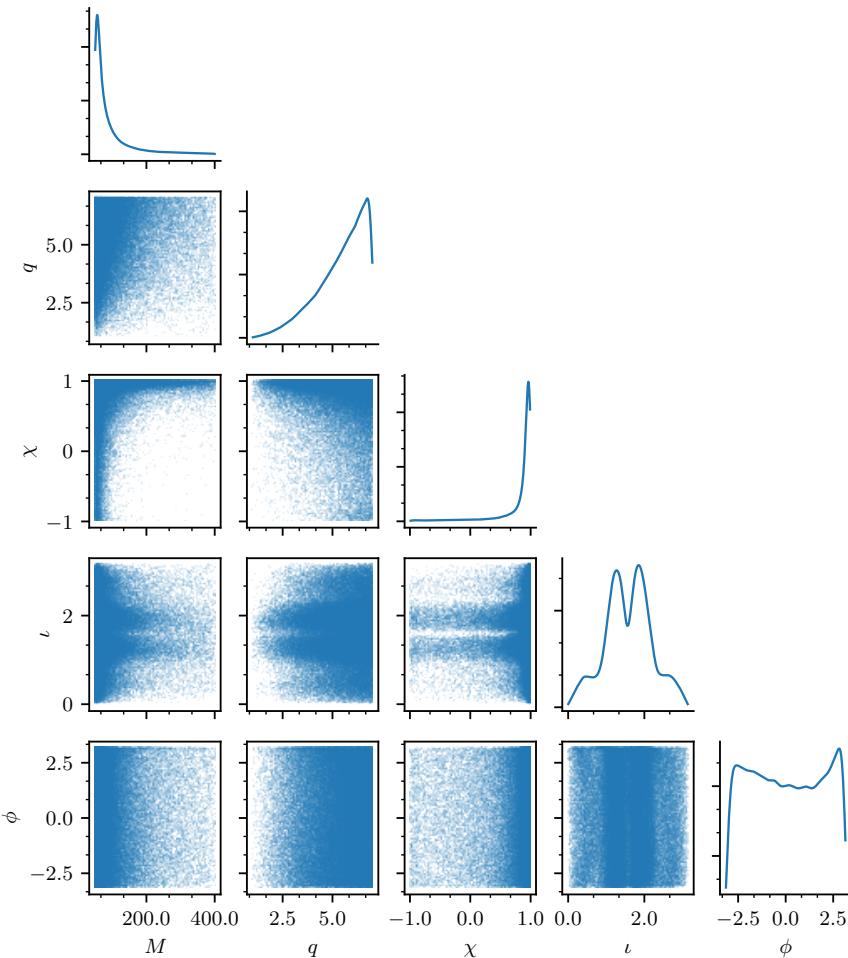


Figure 4.13: Corner plot with the templates of the aligned-spin HM bank described in Sec. 4.5.2. Along the diagonals, we show the histogram of the template number as a function of each coordinate.

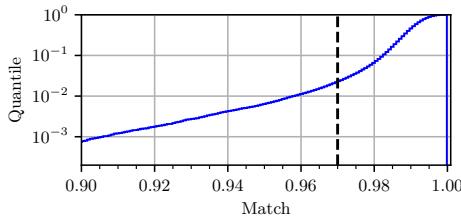


Figure 4.14: Cumulative fitting factor for the aligned-spin HM bank described in Sec. 4.5.2. The histogram is built upon 10^5 injections sampled from the manifold.

have a fitting factor smaller than the target of 0.97, with a median fitting factor of 0.99. We can conclude that the bank provides good coverage of the parameter space. Moreover, the fitting factor is rather constant across all the parameters space. As was also the case for the HM bank introduced in Sec. 4.4.1, there are not regions which are undercovered by the template banks. Also the accuracy of the normalizing flow does not vary too much over the parameter space, showing a bad performance only in the region with high total mass and low mass ratio.

We note that, in order to achieve good performance in the two HM banks presented in this work, we set a covering fraction of only $\eta = 0.8$. This is significantly lower than what we used for the non-HM banks and also lower than the recommended value of $\eta = 0.9$ in [346]. This means that, unlike the non-HM case, the metric match in Eq. 4.5 *underestimates* the “true” match. In this scenario, the covering fraction estimated with the live points (which makes use of the metric) also underestimates the “true” covering fraction. Therefore, a lower value of η is enough to obtain an acceptable coverage. This is not the case for non-HM banks. The reason why this happens only for HM banks is currently not understood and requires more investigation.

4.5.3 Other Possible Applications

The speed of the bank generation, together with the flexibility of the flow in sampling from the parameter space, allows for several novel applications of our work to GW data analysis, besides producing high-dimensional template banks. Without being exhaustive, we discuss below some of the new possibilities.

Selecting the parameter space to target As already discussed, the choice of the parameter space to target in GW searches can be challenging, as it is hard to obtain a reliable forecast of the number of templates needed for accurate coverage. Moving

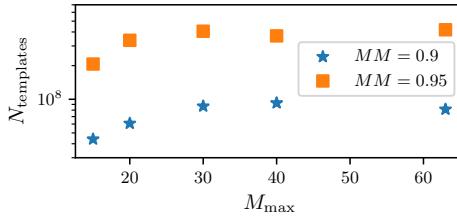


Figure 4.15: Study of the size of a template bank in the neutron star-black hole parameter space. Each point refers to a template bank on the manifold `logMq_s1xz_iota`, covering a total mass range $M \in [M_{\min}, M_{\max}]$. The component masses are limited to $m_1 \in [10, 60] M_{\odot}$ and $m_2 \in [1, 3] M_{\odot}$, with mass ratios $q \in [3.3, 15]$. In the plot we report the number of templates $N_{\text{templates}}$ as a function of the maximum total mass M_{\max} , for different minimal match requirements. The resulting banks are huge, with tens of millions of templates, showing that a search for precessing NSBH binaries is still prohibitively costly.

towards high-dimensional template banks, the number of templates increases by orders of magnitude and the standard stochastic approach suffers from memory issues due to the storage of the waveforms needed for the match calculation. This in turn makes it difficult to even explore high-dimensional spaces, as the current algorithms time-out by the time the bank reaches several million templates

Our method has a low memory footprint and this makes possible to forecast the number of templates in a given parameter space, providing invaluable information to choose an appropriate target for the search. To do so, the interested user might train a normalizing flow on a large region of the parameter space and then place templates in a subregion, without the need to store them. Sampling in a subregion can be easily completed with the use of rejection sampling.

A natural candidate to demonstrate the usefulness of this technique is the precessing NSBH parameter space. Indeed, due to the large mass asymmetry of NSBH systems (i.e. high q), precession has a strong imprint on the waveform, leading to a very large volume to cover by a template bank. To study the number of templates needed to cover the space, we train a normalizing flow model on the manifold `logMq_s1xz_iota`¹¹ for systems with masses $m_1 \in [10, 60] M_{\odot}$ and $m_2 \in [1, 3] M_{\odot}$, with mass ratios $q \in [3.3, 15]$. The other coordinates s_1, θ_1 and ι vary in set $[0, 0.9] \times [-\pi, \pi] \times [0, \pi]$. As above, we use the approximant `IMRPhenomXP`, in a frequency range of [15, 1024] Hz.

To study the parameter space size, we run our template placement algorithm for varying maximum total mass M_{\max} and we measure the number of templates needed

¹¹We neglect any spin on the lighter object, a neutron star. This is physically motivated by the fact that a NS is expected to have low spins [424].

to achieve a covering fraction of $\eta = 0.9$ for different minimal match requirements. Since we do not store and validate the template banks, there is no guarantee that the resulting banks provide a satisfactory coverage. The procedure is just meant to obtain an order of magnitude estimation of the bank size.

As shown in Fig. 4.15, the precessing NSBH parameter space is huge. With a minimal match requirement of 0.9, around 100 million templates are needed to cover the full space. Around half of the templates are in the low total mass region with $M \in [11, 15] M_{\odot}$. The numbers agree with the investigations carried out in [406]. To cover the space with a minimal match of 0.95, around five times more templates are needed.

The magnitude of the precessing NSBH space makes it nearly impossible to use traditional matched-filtering techniques to search for such signals. It thus becomes compelling to either develop new search techniques [406] or to improve the computational power available.

Thanks to our method, similar estimates can easily be done for other regions of the BBH parameter space (e.g. targeting eccentric BBHs), thus providing invaluable information to plan future high-dimensional GW searches.

Manifold selection The metric eigenvalues and eigenvectors can give an interesting piece of information about the relative importance of the coordinates of the manifold. Let λ_i and \mathbf{v}_i be the i -th eigenvalue and eigenvector respectively of the metric M_{ij} . We can think of each eigenvector λ_i as a measure of the relative importance of the eigenvector \mathbf{v}_i , which represents a linear combination of the coordinates. We can then introduce the following quantity for each coordinate j , which we call *coordinate importance*:

$$\mathcal{I}_j = \left| \sum_i \lambda_i (\mathbf{v}_i)_j \right| \quad (4.27)$$

where $(\mathbf{v}_i)_j$ is the j -th component of the i -th eigenvector. It is a weighted average over the projection of each eigenvector along a given coordinate. Heuristically, an “important” coordinate will give a larger contribution to the “important” eigenvectors (i.e. with larger eigenvalues).

This quantity might be used to create a hierarchy among the coordinates and, when choosing the manifold to cover, it can offer a useful criteria to decide which quantities to include in the bank. For example, in the manifold `10Mq_s1xyz_s2z_-iota_phi`, the variable $\log_{10} M$ has an importance of 5×10^4 , while variables q, s_1, θ_1 and s_{2z} have importance two orders of magnitude less. This implies that a template bank must include (besides the total mass) all the variables q, s_1, θ_1 and s_{2z} . On the other hand, coordinates φ_1 (controlling the magnitude of s_{1y}) and the angles ι and

ϕ have an “importance” of one order of magnitude less than all the other quantities. As a consequence, the latter three play a smaller role in covering the space and they can be possibly ignored (or perhaps only one of them can be included).

Of course, this line of reasoning is heuristic and whether a manifold is suitable or not to cover the space must be checked by means of an injection study. However, the study of the relative importance between coordinates can give an educated guess on the manifold to cover and serve as a starting point for the trial and error process of manifold selection.

A proposal for the stochastic template placement Our normalizing flow finds an obvious application within a stochastic placement algorithm. According to the stochastic algorithm, template proposals are randomly drawn from an analytical PDF, which is specifically design to approximate Eq. (4.10) in the non-spinning case. A good proposal is crucial to reduce the template rejection rate, hence reducing the overall run time.

The normalizing flow is a natural candidate for a proposal distribution, since it goes beyond the non-spinning BBH approximation, allowing for more physics to be considered. Implementing a normalizing flow within the stochastic algorithm will most likely provide a computational benefit, due to a more efficient proposal.

Generating datasets for machine learning applications The recent years have seen a burst of machine learning application to GW data analysis, covering all fields of the analysis of compact binary systems from waveform modelling [5, 6, 270, 411] to GW searches [428–431] and parameter estimation [288–290, 292, 432].

For all these applications, it is crucial to have high quality datasets of waveforms for training purposes. The goodness and the applicability of the model strongly relies on the distribution of waveforms in the dataset and substantial time is often spent in tuning the dataset composition to achieve optimal performance. The waveforms in such datasets can be sampled using our normalizing flow model, thus covering the space accurately. In many cases this may prove beneficial.

4.6 Future Prospects

Clearly, our work can be improved and expanded in several directions. In this section, we discuss some possible advancements.

Introducing a new metric As shown in App. 4.B, the Hessian of the match (with which we identify the metric) does not always approximate the behaviour of the true

match in a neighborhood of a point. For instance, on the manifold `Mq_s1xyz`, consider the ellipse \mathcal{E}_0 , centered on $\theta_0 = (10 M_\odot, 7, 0.6, 2, 2)$ of all the points θ with *metric* match with the center higher than 0.97. It turns out that only $\sim 50\%$ of the points inside \mathcal{E}_0 have a match higher than 0.97. The situation gets worse for smaller mass ratio, when the metric determinant vanishes, and it can significantly vary among different manifolds.

While this hasn't affected (too much) the effectualness of our template bank, the failure of the metric approximation is concerning and can negatively influence the placement, especially in presence of a parameter with a small impact on the waveform. The interested reader is encouraged to read App. 4.B.

Exploring different flow architectures In this work, we only considered Masked Autoregressive Layers for our normalizing flow architecture. Of course, other choices are available in the literature and could possibly improve the flow accuracy. Further work should implement some of these and assess the (possible) gain in accuracy. Possible transformations include coupling layers [433, 434] or residual flows [435, 436].

As discussed in Sec. 4.2.2, it is very beneficial to use a transformation like Eq. (4.14) as the first layer of the normalizing flow. Future work can find a different transformation offering better performance.

Estimating the covering fraction with importance sampling An accurate evaluation of the covering fraction in Eq. (4.20) is crucial to providing a realistic estimation of the template number and hence good coverage. Currently we estimate the covering fraction by using the approximation to the volume element given by the normalizing flow. We can increase the accuracy by computing the integral in Eq. (4.20) with importance sampling:

$$\hat{\eta}(\mathcal{T}) \simeq \frac{1}{\sum_i w_i} \sum_i c(\theta_i) w_i \quad (4.28)$$

where the live points are sampled from the flow and are weighted with weights $w_i = \frac{\sqrt{|M(\theta_i)|}}{p^{\text{flow}}(\theta_i)}$. The weights make sure that we evaluate the unapproximated version of the integral, i.e. using the true volume element and not its flow approximation.

In a practical application, it is wise to prevent the weights to grow indefinitely, as this can negatively impact the estimation of the covering fraction. For this reason, we clip the weights to a maximum value of W : $w_i = \min\left(\frac{\sqrt{|M(\theta_i)|}}{p^{\text{flow}}(\theta_i)}, W_{\max}\right)$. The tuning of W_{\max} deserves more attention, as it can really impact the bank performance.

Some tests have shown that importance sampling delivers larger banks, thus with better coverage but with an increased variance in the number of templates. However,

in some occasions, one or a few live points can dominate the sum (i.e. have very large weight), making the covering fraction computation less robust in case of flow inaccuracies. More work is required to treat such cases and successfully implement this new feature.

Exploring different placement methods While the random template placement method in use has proven its effectiveness, other alternatives are certainly possible. A different placement method is appealing to reduce the bank size without degrading its performance, as random template banks tend to place more templates than needed.

First, one could use the metric to reject templates that are too close to each other. This would be a variation of the stochastic algorithm, where distances are computed with the metric and not with the true match. While this may prove unfeasibly slow in some cases, it can still be computationally more efficient than with the brute force match computation. As a compromise, a random template bank with low covering fraction and minimal match might be given as starting point for the iteration (i.e. a seed bank).

One could also devise alternative strategies to sample from the flow latent space, such as using quasi Monte Carlo sampling or even setting points on a lattice. Since the coordinates of the templates will be correlated with each other, we cannot compute iteratively the covering fraction as described in Sec. 4.2.3. For this reason the suitable bank size needs to be computed with other methods, before selecting the templates.

Regardless of the placement method, the templates in a bank may still not be placed optimally, creating over(under)-dense regions. This is especially true for the *random* method used here. For this reason, it may be beneficial to add a post-processing step to move or remove some templates [437].

Encoding the metric into the flow? A fascinating path to explore is to encode information about the metric M_{ij} inside the flow transformation. So far, the normalizing flow Φ_W is trained in such a way that the determinant of the Jacobian $\det J_{\Phi_W}$ matches the determinant of the metric. Thus, among the $\frac{D(D-1)}{2}$ free components of J_{Φ_W} , only one of them is constrained during the training. This leaves a lot of degeneracy in J_{Φ_W} . One could break such degeneracy by imposing the additional constraint that the Jacobian of the flow matches the metric M_{ij} :

$$(J_{\Phi_W})_{ij} \simeq M_{ij}. \quad (4.29)$$

Such constraint should be imposed by introducing a suitable loss function. The approach would involve a much harder optimization problem and it remains to be assessed whether the flow has enough representation power to solve such problem.

A flow trained in this way would create an isometry (i.e. distance preserving transformation) between the latent space and the physical space. According to differential geometry, this is not possible, unless the M_{ij} has zero curvature, which is not the case in general. A possible way out could be to embed the manifold of signals in an higher dimensional flat manifold, which would guarantee the existence of a solution.

As outlined, there are many open questions and issues to solve, which require significant work. The reward however would be significant: the flow would parameterize a distance preserving (and not only volume preserving) transformation, which can be used for high dimensional fast stochastic placement or even geometric placement – the holy grail of bank generation.

4.7 Final Remarks

In this chapter, we presented a novel method to generate template banks covering a high-dimensional manifold of (possibly) precessing and/or HM and/or eccentric BBH signals. Key to our method is the metric M_{ij} and the derived volume element $\sqrt{|M|}$. The latter defines the number of templates that should cover an infinitesimal volume and can be seen as a probability measure on the space. We derive here for the first time an expression for the metric suitable for precessing and/or HM signals (see App. 4.A). The metric is written in terms of the gradients of the waveform polarizations and is numerically stable.

To sample the templates, we introduce a novel normalizing flow model, which serves the twofold purpose of sampling from the space and providing a fast-to-compute approximation to $\sqrt{|M|}$. Once we are able to sample from the space, we place templates using the *random* algorithm, which is fast and suitable to cover high-dimensional spaces. This comes at the price of a larger bank than would be produced with the state-of-the-art stochastic algorithm, although the over-coverage becomes less severe as the number of dimensions, and correspondingly the overall size of the bank, increases.

We validate our code by evaluating the normalizing flow accuracy and the robustness of the random placement. Moreover, with a few hours of computation, we were able to reproduce two template banks existing in the literature obtained with independent codes - a non-spinning HM bank [328] and an aligned-spin bank [380].

To demonstrate the capabilities of our code, we generate two large template banks covering systems for which no or little searches have been performed: a precessing bank gathering 1.6 million templates (Sec. 4.5.1) and an aligned-spin HM bank formed by 2.1 million templates (Sec. 4.5.2). We show that the two banks satisfac-

torily cover the space. They were both produced in a matter of hours, with minimal CPU and memory usage. We also discuss other possible applications of our method, including the optimization of the template proposal of the stochastic algorithm, the selection of a suitable parameter space for a GW search and the generation of datasets of waveforms for the training of machine learning models. Our code is publicly available as a package `mbank` [410] and comes with a large number of tools to simplify the bank generation and validation.

As a final remark, we stress that our work will enable the GW community to run searches on novel regions of the BBH parameter space. Being able to generate a high dimensional bank in a few hours, the computational cost of searching new regions of the parameter space will be dominated by the actual cost of the analysis rather than the cost of prior steps. This will allow for optimal resource allocation to search for signatures of precession, eccentricity and/or HMs, hopefully leading to exciting physics discoveries. For this reason, we will heavily rely on `mbank` to generate the template bank needed for a precessing search. Its speed and flexibility quickly allowed us to explore different regions of the parameter space, allowing us to identify a suitable search target and to cover it with templates.

4.A Details of the Metric Computation

In this Appendix we report the details of the derivation of Eq. (4.6), as well as the computation of the Hessian H of the overlap in Eq. (4.1) in terms of the gradients of the waveform $h(\theta)$. In what follows, we define $(h_1|h_2)$ and $[h_1|h_2]$ to be the real and imaginary part, respectively, of $\langle h_1|h_2 \rangle$.

We begin by expanding the quantity $\mathcal{M}(\theta + \Delta\theta, \theta)$ for $\Delta\theta$ around 0. Since $\mathcal{M}(\theta + \Delta\theta, \theta)$ has a maximum for $\Delta\theta = 0$, the leading term is quadratic in $\Delta\theta$. We obtain:

$$\begin{aligned} \mathcal{M}(\theta + \Delta\theta, \theta) &= \max_{\Delta t} \mathcal{O}(\theta + \Delta\theta, \theta, \Delta t) \\ &= \max_{\Delta t} \left\{ 1 + \frac{1}{2} [\partial_{ij} \mathcal{O} \Delta\theta_i \Delta\theta_j + 2\partial_{it} \mathcal{O} \Delta\theta_i \Delta t + \partial_{tt} \mathcal{O} (\Delta t)^2] \right\} \\ &= 1 + \frac{1}{2} \left[\partial_{ij} \mathcal{O} - \frac{\partial_{it} \mathcal{O} \partial_{jt} \mathcal{O}}{\partial_{tt} \mathcal{O}} \right] \Delta\theta_i \Delta\theta_j \end{aligned} \quad (4.30)$$

where all the derivatives are evaluated at $\Delta\theta = \Delta t = 0$ and the explicit time maximization yields $\Delta t = -\frac{\partial_{it} \mathcal{O} \Delta\theta_i}{\partial_{tt} \mathcal{O}}$.

From Eq. (4.30), we can read the expression for the metric in Eq. (4.6) recognizing in the derivatives $\partial\partial\mathcal{O}|_{\Delta\theta, \Delta t=0}$ the components of the Hessian matrix H of the overlap.

We now compute the Hessian H of the overlap in terms of the gradients of the *normalized* waveforms. For notational convenience, we set $h_+(\theta_1)e^{ift} = s$, we drop

any dependence on θ_2 and we understand $\mu = i, t$. We have:

$$\begin{aligned}\partial_\mu \mathcal{O} = \frac{1}{\mathcal{O}} \frac{1}{1 - \hat{h}_{+ \times}^2} & \left[(\partial_\mu \hat{s} | \hat{h}_+) (\hat{s} | \hat{h}_+) + (\partial_\mu \hat{s} | \hat{h}_\times) (\hat{s} | \hat{h}_\times) \right. \\ & \left. - (\partial_\mu \hat{s} | \hat{h}_+) (\hat{s} | \hat{h}_\times) h_{+ \times} - (\partial_\mu \hat{s} | \hat{h}_\times) (\hat{s} | \hat{h}_+) h_{+ \times} \right]\end{aligned}\quad (4.31)$$

Differentiating another time, after some rearrangements, we get:

$$H_{tt} = -(\hat{h}_+ | \hat{h}_+ f^2) + \frac{1}{1 - \hat{h}_{+ \times}^2} [\hat{h}_\times | \hat{h}_+ f]^2 \quad (4.32)$$

$$H_{ti} = [\hat{h}_+ | \partial_i \hat{h}_+ f] - \frac{1}{1 - \hat{h}_{+ \times}^2} (\hat{h}_\times | \partial_i \hat{h}_+) [\hat{h}_\times | \hat{h}_+ f] \quad (4.33)$$

$$H_{ij} = (\hat{h}_+ | \partial_i \partial_j \hat{h}_+) + \frac{1}{1 - \hat{h}_{+ \times}^2} (\hat{h}_\times | \partial_i \hat{h}_+) (\hat{h}_\times | \partial_j \hat{h}_+) \quad (4.34)$$

To move further, we express the normalized waveform derivatives in terms of the un-normalized ones:

- $\partial_i \langle h | h \rangle = \langle \partial_i h | h \rangle + \langle h | \partial_i h \rangle = 2(h | \partial_i h)$
- $\partial_i \hat{h} = \frac{1}{(h | h)^{3/2}} [(h | h) \partial_i h - (h | \partial_i h) h]$
- $\partial_t \hat{h} = i f \hat{h} = i f \frac{h}{(h | h)^{1/2}}$
- $\partial_i \partial_j \hat{h} = \frac{1}{(h | h)^{1/2}} \partial_{ij} h + 3 \frac{1}{(h | h)^{5/2}} (h | \partial_i h) (h | \partial_j h) h$
 $- \frac{1}{(h | h)^{3/2}} [(h | \partial_{ij} h) h + (\partial_i h | \partial_j h) h + 2(h | \partial_{(i} h) \partial_{j)} h]$

where $A_{(ij)} = \frac{1}{2}(A_{ij} + A_{ji})$ denotes symmetrization.

Plugging this into the equations (4.32)-(4.34), we get:

$$H_{tt} = -\frac{1}{h_{++}} (h_+ | f^2 h_+) + \frac{1}{1 - \hat{h}_{+ \times}^2} \frac{1}{h_{++} h_{\times \times}} [h_\times | f h_+]^2 \quad (4.35)$$

$$\begin{aligned}H_{ti} = -\frac{1}{h_{++}} (h_+ | f \partial_i h_+) - \frac{1}{1 - \hat{h}_{+ \times}^2} \frac{1}{h_{++} h_{\times \times}} [h_\times | f h_+] (h_\times | \partial_i h_+) \\ + \frac{\hat{h}_{+ \times}}{1 - \hat{h}_{+ \times}^2} \frac{1}{h_{++}^{3/2} h_{\times \times}^{1/2}} [h_\times | f h_+] (h_+ | \partial_i h_+)\end{aligned}\quad (4.36)$$

$$\begin{aligned}H_{ij} = -\frac{1}{h_{++}} (\partial_i h_+ | \partial_j h_+) + \frac{1}{1 - \hat{h}_{+ \times}^2} \frac{1}{h_{++}^2} (h_+ | \partial_i h_+) (h_+ | \partial_j h_+) \\ + \frac{1}{1 - \hat{h}_{+ \times}^2} \frac{1}{h_{++} h_{\times \times}} (h_\times | \partial_i h_+) (h_\times | \partial_j h_+) \\ - \frac{2 \hat{h}_{+ \times}}{1 - \hat{h}_{+ \times}^2} \frac{1}{h_{++}^{3/2} h_{\times \times}^{1/2}} (h_\times | \partial_{(i} h_+) (h_+ | \partial_{j)} h_+)\end{aligned}\quad (4.37)$$

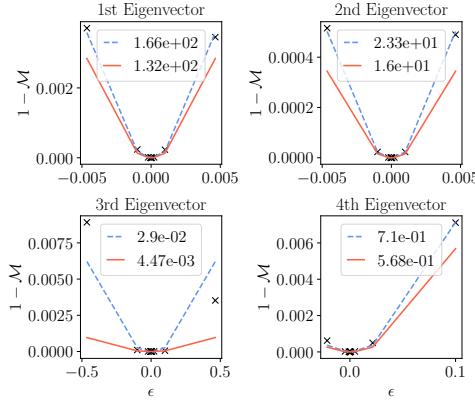


Figure 4.16: For each eigenvector of the metric, we compute the empirical relation between the mis-match $1 - \mathcal{M}$ and the distance ϵ of points along the eigenvector direction. The solid line shows the relation predicted by the metric, while the dashed line shows a parabolic fit. In the legend are reported the quadratic coefficients of both lines.

where we defined $h_{\cdot*} = (h_{\cdot} | h_{\cdot*})$.

These expressions, together with Eq. (4.6) fully specify the metric. The gradients $\partial_i h$ of the waveform can be computed with a finite difference scheme or analytically for a number of surrogate waveform models [5, 6, 270, 411].

The non precessing limit can be recovered by setting $h_{\times} = ih_{+}$ and $h_{+\times} = 0$:

$$H_{tt} = \frac{1}{h_{++}^2} (h_{+} | f h_{+})^2 - \frac{1}{h_{++}^2} (h_{+} | f^2 h_{+}) \quad (4.38)$$

$$H_{ti} = \frac{1}{h_{++}^2} [h_{+} | \partial_i h_{+}] (h_{+} | h_{+} f) - \frac{1}{h_{++}^2} [h_{+} | f \partial_i h_{+}] \quad (4.39)$$

$$H_{ij} = \frac{1}{h_{++}^2} \left\{ (h_{+} | \partial_i h_{+}) (h_{+} | \partial_j h_{+}) + [h_{+} | \partial_i h_{+}] [h_{+} | \partial_j h_{+}] \right\} - \frac{1}{h_{++}^2} (\partial_i h_{+} | \partial_j h_{+}) \quad (4.40)$$

4.B Alternative Definitions for the Metric

Throughout this paper, we identified the metric with the Hessian of the overlap (see Eq. (4.6)). While this is widely used in the literature [336, 345] and has been proven to provide reliable template banks, it still has some undesirable properties. To show this, we compute the metric at point $\theta_0 = (20 M_{\odot}, 3., 0.7, 1.8)$ of manifold `Mq_s1xz`, described in Sec. 4.3.1, and we compute its eigenvalues $\alpha^{(i)}$ and eigenvectors $v^{(i)}$. We then compute the match $\mathcal{M}_{\epsilon}^{(i)}$ between θ_0 and the point $\theta_{\epsilon}^{(i)} = \theta_0 + \epsilon v^{(i)}$, located at a

distance ϵ along i -th eigenvector. Finally, we compute the coefficient α of the Taylor expansion $1 - \mathcal{M}_\epsilon^{(i)} = \alpha\epsilon^2$. α corresponds to the i -th eigenvalue and in principle, it should be close to its value.

In Fig. 4.16, we plot the fitted relation between $1 - \mathcal{M}$ and ϵ for each eigenvector, as well as the one computed with the metric. In the legend we report the α coefficient (dashed blue line) and the eigenvalue of the metric (solid orange line). The striking feature we note in Fig. 4.16 is that the eigenvalue is consistently smaller than the fitted α coefficient, sometimes by an order of magnitude. This means that the Hessian, which is computed for $\epsilon \rightarrow 0$, is not able to extrapolate the behaviour of $1 - \mathcal{M}(\epsilon)$ even at modestly large value of ϵ : the metric approximation to the match loses its predictivity as a measure of distance. The problem becomes more severe in high-dimensional manifolds. On the other hand, since the banks generated with the Hessian metric show nice coverage, one may argue that the *volume* estimate provided by the Hessian is still accurate enough for our purposes.

As a way out, we could redefine the matrix $M_{ij}(\theta)$ to a more suitable expression, departing from the Hessian. The goodness of the metric expression may depend on the application and on the range of validity of the approximation. The tensor field $M_{ij}(\theta)$ can be computed through an optimization problem, where we minimize the discrepancy between the two quantities in Eq. (4.5), encoded into a *loss function*. The loss function depends on the values of the matrix elements M'_{ij} :

$$\mathcal{L}_\theta(M'_{ij}) = \int_{\{d(\theta, \theta') < d_{\text{target}}\}} d^D \theta' [1 - \mathcal{M}(\theta, \theta') - M'_{ij} \Delta \theta_i \Delta \theta_j]^2 \quad (4.41)$$

where the integration extends on a D -ball with radius d_{target} centered around θ and d_{target} is a tunable parameter, which controls the validity of the approximation.

At any given point θ , the components $M_{ij}(\theta)$ of the metric are selected by minimizing the above loss:

$$M_{ij}(\theta) = \arg \min_{M'_{ij}} \mathcal{L}_\theta(M'_{ij}). \quad (4.42)$$

Although the minimization can be tackled with standard techniques, it requires many evaluations of Eq. (4.4) and the ability to sample from a “complex” set such as $\{d(\theta, \theta') < d_{\text{target}}\}$.

While in most cases this may prove unfeasible, future work could solve the problem in Eq. (4.42) at a manageable cost. This may be beneficial to many data analysis applications, such as template placement and Fisher information matrix studies. A number of alternative metric expressions, coming from different heuristic optimization strategies, are already available in `mbank`, although not fully validated.

4.C Computing the Volume of the Parameter Space

Since the number of templates is proportional to the volume of the parameter space [336], estimating this volume can be valuable for forecasting the size of a template bank. The estimation can be easily done by *importance sampling*, using samples from the normalizing flow, since the latter is trained to reproduce the volume element.

The volume of the parameter space \mathcal{B}_D is defined as:

$$\mathcal{V} = \int_{\mathcal{B}_D} d^D \theta \sqrt{\det M(\theta)} \quad (4.43)$$

$$= \int_{S_{\text{flow}}} d^D \theta \sqrt{\det M(\theta)} \mathcal{I}_{\mathcal{B}_D}(\theta) \quad (4.44)$$

where in the last equality we compute the integral on the support of the flow $S_{\text{flow}} \supseteq \mathcal{B}_D$ and we introduced the indicator function $\mathcal{I}_{\mathcal{B}_D}$ which is non-zero only on the manifold \mathcal{B}_D . Eq. (4.44) can be numerically evaluated by importance sampling:

$$\mathcal{V} \simeq \frac{1}{N} \sum_i \frac{\sqrt{\det M(\theta_i)}}{p^{\text{flow}}(\theta_i)} \mathcal{I}_{\mathcal{B}_D}(\theta_i) \quad (4.45)$$

with $\theta_i \sim p^{\text{flow}}$. The normalizing flow ensures a low variance in the volume estimation.

Eq. (4.45) involves several metric evaluations, which has some computational cost. To further reduce the computational cost, we can use the fact that, after the training procedure, the flow approximates the volume element as follows:

$$\log p^{\text{flow}} - \log \sqrt{|M|} + C \simeq 0 \quad (4.46)$$

where C is the trainable constant appearing in Eq. (4.16). Hence we can replace $\frac{\sqrt{\det M(\theta_i)}}{p^{\text{flow}}(\theta_i)}$ in Eq. (4.45) simply with e^C . The volume estimation is then reduced to computing the fraction of the volume of S_{flow} covered by \mathcal{B}_D :

$$\mathcal{V} \simeq e^C \frac{1}{N} \sum_i \mathcal{I}_{\mathcal{B}_D}(\theta_i) \quad (4.47)$$

where again $\theta_i \sim p^{\text{flow}}$. The goodness of such approximation is closely related to the flow performance, as studied in Sec. 4.3.1 (see also Fig. 4.1).

Once an estimation of the volume is available, the number of templates can be obtained by noting [336] that in a lattice, given a minimal match MM , the average spacing d between template is:

$$d(MM) = 2 \sqrt{\frac{1 - MM}{D}}. \quad (4.48)$$

Hence, roughly speaking, the number of templates N needed to cover the volume \mathcal{V} is given by:

$$N = \frac{\mathcal{V}}{d(MM)^D}. \quad (4.49)$$

CHAPTER 5

Generalized Signal Consistency Test

5.1 Introduction

As discussed in Sec. 3.4.2, in the presence of non-Gaussian transient noise bursts [211, 213–216] a matched-filter routine can occasionally output triggers with large SNR values. As such triggers can contaminate the background of a search, hence lowering its sensitivity, several statistical tests, known as signal consistency tests [77,79,82,382–387,389,391], have been designed to improve the separation between GW candidates and false alarms caused by the detectors’ non-Gaussian noise.

Expanding our searches to include signals from precessing binaries [66,74,327, 405, 406, 438] and/or binaries with higher-order modes (HMs) [328, 373, 407–409] necessitates the generalization of those signal consistency tests to a more versatile framework. Neglecting to update these tests can result in decreased search sensitivity, potentially offsetting the benefits of using a more diverse set of templates. While the χ^2 time-frequency signal consistency test [382] and its variant, the sine-Gaussian χ^2 discriminator [384], have been successfully applied in searches including higher-order modes [328, 373] and precessing signals [406], little attention has been given to generalizing other types of signal consistency tests, or even to the development of new ones.

In particular, in this work based on [2] we focus on the autocorrelation-based least-squares test Eq. (3.50), denoted ξ^2 , which is currently utilized by the GstLAL search pipeline. The traditional test relies on the assumption that the signal to detect is a circular aligned-spin binary system where no HMs are considered. Indeed, as

detailed in Sec. 1.3.3 in the discussion around Eq. (1.101), in this scenario the system is symmetric under reflection across the orbital plane and this equatorial symmetry translates into a simple relation between the two polarizations of the GW emitted

$$\tilde{h}_+ \propto i \tilde{h}_\times, \quad (5.1)$$

which allows to obtain a convenient and computationally efficient expression for the test. In this Chapter, we describe how the simplicity of Eq. (5.1) enters the expression for the currently used ξ^2 signal consistency test [77] and we describe how to move away from the “aligned-spin and no HM” assumption. Consequently, we introduce a new signal consistency test that does not make *any assumption* about the nature of the signal to detect. Thus, while the test is primarily motivated to search for precessing and/or HM signals, it can be applied to a matched-filtering search for any type of gravitational-wave signal. In Sec. 5.2, we discuss in details the “standard” test and show how it can be generalized to precessing and/or HM binaries, by introducing our new generalized signal consistency test ξ_{sym}^2 . Moreover, we provide a computationally convenient approximate expression ξ_{mix}^2 to the novel test. In Sec. 5.3 we discuss the performance of the newly introduced test and its approximated version. Sec. 5.4 gathers some final remarks and future prospects.

5.2 A New Generalized ξ^2 Signal consistency Test

As introduced in Sec. 3.2.3, the generic output of the matched-filtering procedure is the complex timeseries $z(t)$:

$$z(t) = (d|h_R)(t) + i(d|h_I)(t) \quad (5.2)$$

where h_R and h_I are two *normalized* real templates. The *real* filters h_R, h_I are chosen to maximise the SNR timeseries $\rho(t) = |z(t)|$ at the time where a GW signal is present in the data and their expression depends on the assumptions about the nature of the GW signal to search.

Given a trigger at time $t = 0$, the ξ^2 signal consistency test relies on predicting the SNR timeseries $z(t)$ obtained by filtering a signal h with a matching templates. The predicted timeseries $R(t)$ is then compared to the measured timeseries $z(t)$ to compute the squared residual timeseries:

$$\xi^2(t) = |z(t) - R(t)|^2. \quad (5.3)$$

We can integrate the residual timeseries to obtain the ξ^2 statistic:

$$\xi^2 = \frac{\int_{-\delta t}^{\delta t} dt |z(t) - R(t)|^2}{\int_{-\delta t}^{\delta t} dt E[\xi^2(t)]} \quad (5.4)$$

where the integral extends on a short time window $[-\delta t, \delta t]$ around the trigger time. To obtain the ξ^2 statistic, the integral of the residual timeseries is normalized by the integral of the expected value $E[\xi^2(t)]$ over different Gaussian noise realizations of the residual timeseries without a signal buried in the noise. Clearly, its value depends on the template considered. It is convenient to express δt in terms of the so-called autocorrelation length (ACL) [77], defined as the number of samples in the time-window $[-\delta t, \delta t]$ at a given sample rate f_{sampling} , so that $\delta t = (\text{ACL} - 1)/(2 f_{\text{sampling}})$.

The ξ^2 defined above can be used by the GW search pipelines to veto some loud triggers. If a trigger is caused by a noise fluctuation or non-Gaussian noise transient bursts, the discrepancy between the expected and measured SNR timeseries will be large, leading to a large value of ξ^2 . This can be used to downrank certain triggers, with large improvement in sensitivity.

5.2.1 The “Standard” Signal Consistency Test

In the case of aligned-spin systems, where no HMs are considered, the two templates h_R, h_I used to filter the data are given by \hat{h}_p, \hat{h}_c introduced in Eq. (3.21) and the matched-filtering procedure outputs the “standard” SNR timeries $z_{\text{std}}(t)$ Eq. (3.26). Thanks to the equatorial symmetry of the “standard” systems, mathematically expressed by Eq. (5.1), (i) the two templates are related by the simple expression $\tilde{h}_R = i\tilde{h}_I$ and (ii) the observed signal at the detector Eq. (3.23) in frequency domain is also proportional to \tilde{h}_p . Thanks to this remarkable property of the signal, the predicted SNR timeseries $R_{\text{std}}(t)$ around a trigger at $t = 0$ is given by:

$$\begin{aligned} R_{\text{std}}(t) &= (h|\hat{h}_p)(t) + i(h|\hat{h}_c)(t) \\ &= z(0) \left\{ (\hat{h}_p|\hat{h}_p)(t) + i(\hat{h}_p|\hat{h}_c)(t) \right\} \end{aligned} \quad (5.5)$$

where, to simplify the expression above, we used the fact that $h = \frac{1}{D_{\text{eff}}} e^{i\phi} \hat{h}_p$ and that $z(0) = \frac{1}{D_{\text{eff}}} e^{-i\phi}$, since $(\hat{h}_c|\hat{h}_p)(t = 0)$ is equal to zero. By comparing the measured timeseries $z_{\text{std}}(t)$ and the expected timeseries $R_{\text{std}}(t)$ with Eq. (5.4), one can compute ξ_{std}^2 for a non-precessing search. It is a simple exercise to compute the expected value of the residual timeseries over different noise realizations [77, App. A]:

$$E[\xi_{\text{std}}^2(t)] = 2 - 2 \left| (\hat{h}_p|\hat{h}_p)(t) + i(\hat{h}_p|\hat{h}_c)(t) \right|^2 \quad (5.6)$$

The integral of this expression can be used as a normalization for the ξ^2 statistic¹.

¹Note that the normalization of ξ^2 computed in Eq. (5.6) is equivalent to the denominator of Eq. (3.50) introduced in Sec. 3.4.2.

By using again $\tilde{h}_p = i\hat{h}_c$ and the identity $[a|ib] = (a|b)$, the “standard” SNR Eq. (3.26) can also be expressed in terms of \hat{h}_p only:

$$\rho_{\text{std}}(t) = \sqrt{(h|\hat{h}_p)^2 + [h|\hat{h}_p]^2}. \quad (5.7)$$

Of course the same is readily done for the $R_{\text{std}}(t)$:

$$\begin{aligned} R_{\text{std}}(t) &= z(0) \left\{ (\hat{h}_p|\hat{h}_p)(t) + i[\hat{h}_p|\hat{h}_p](t) \right\} \\ &= z(0) \langle \hat{h}_p | \hat{h}_p \rangle(t) \end{aligned} \quad (5.8)$$

The quantity $\langle \hat{h}_p | \hat{h}_p \rangle(t)$ is sometimes called template autocorrelation. For “standard” signals, the Eqs. (5.5) and (5.8) are equivalent. However, for precessing and/or HM they may give very different results, due to the breaking of the symmetry Eq. (3.22).

We close by noting that the predicted timeseries is given by a product of a trigger-dependent scalar and a complex template-dependent timeseries

$$R(t) = \text{Trigger} \times \text{Complex timeseries}$$

This arises directly from the fact that the signal model Eq. (3.23) presents the same convenient factorization. Such factorization makes the ξ^2 evaluation particularly computationally convenient and hence the ξ^2 test computationally attractive. As we will see, the use of precessing and/or HM templates breaks this factorization, as Eq. (3.23) is no longer valid.

5.2.2 The Novel “Symphony” Signal Consistency Test

The simplicity of the “standard” case arises directly from the symmetry between the two polarizations $\tilde{h}_+ = i\hat{h}_\times$, which allows us to conveniently factorize the signal model and the predicted SNR timeseries. However, in general this is not possible and a different expression for the filters and the expected SNR timeseries needs to be computed. As we saw in Sec. 3.2.2, with the *only* assumptions that the observed signal is a linear combination of the two polarizations, the appropriate filters for the interferometric data are given by [183, 328, 427]:

$$\begin{aligned} h_R &= \hat{h}_+ \\ h_I &= \hat{h}_\perp = \frac{1}{\sqrt{1 - \hat{h}_{+ \times}^2}} (\hat{h}_\times - \hat{h}_{+ \times} \hat{h}_+) \end{aligned} \quad (5.9)$$

where the real quantity $\hat{h}_{+ \times} = (\hat{h}_+ | \hat{h}_\times)(t = 0)$ is a crucial measure of the precession and/or HM content of a template. The non-precessing non-HM limit can be recovered

by $\hat{h}_{+\times} = 0$. We call *symphony SNR* the SNR timeseries produced with the templates above:

$$z_{\text{sym}}(t) = (h|\hat{h}_+) + i(h|\hat{h}_\perp). \quad (5.10)$$

Note that as in the “standard” case, the two templates \hat{h}_+ and \hat{h}_\perp are orthogonal vectors, i.e. $(\hat{h}_+|\hat{h}_\perp) = 0$. Indeed, the vectors for \hat{h}_+ , \hat{h}_\perp follows the Gram-Schmidt “orthogonalization” prescription to create a set of orthonormal basis from the set of basis vectors $\{\hat{h}_+, \hat{h}_\times\}$. Obviously, the fact that the two filters are orthogonal doesn’t mean that in general \hat{h}_+ and \hat{h}_\perp are related by a simple expression such as Eq. (3.22).

To compute the predicted timeseries $R_{\text{sym}}(t)$, it is convenient to rewrite the signal model Eq. (2.6) in terms of \hat{h}_+ and \hat{h}_\perp

$$h = \mathcal{A}_+ \hat{h}_+ + \mathcal{A}_\perp \hat{h}_\perp, \quad (5.11)$$

where we absorbed into \mathcal{A}_+ , \mathcal{A}_\times an overall scaling factor (which depends on the source distance and on the sky location, as well as on $\hat{h}_{+\times}$). With this definition, the predicted SNR timeseries is given by:

$$\begin{aligned} R_{\text{sym}}(t) &= (h|\hat{h}_+)(t) + i(h|\hat{h}_\perp)(t) \\ &= \mathcal{A}_+ \hat{h}_{++}(t) + \mathcal{A}_\perp \hat{h}_{\perp+}(t) \\ &\quad + i\mathcal{A}_+ \hat{h}_{+\perp}(t) + i\mathcal{A}_\perp \hat{h}_{\perp\perp}(t) \end{aligned} \quad (5.12)$$

where to shorten the notation we defined

$$\hat{h}_{\bullet\star}(t) = (\hat{h}_\bullet|\hat{h}_\star)(t) \quad \text{with } \bullet, \star = +, \times, \perp \quad (5.13)$$

and we identify $\mathcal{A}_{+/\perp}$ with the real and imaginary part respectively of the trigger $z(0) = \mathcal{A}_+ + i\mathcal{A}_\perp$. This directly arises from the setting $R_{\text{sym}}(0) = z(0)$ and recognizing that $\hat{h}_{\perp+}(0) = \hat{h}_{+\perp}(0)$ by definition.

The expression for $R_{\text{sym}}(t)$ Eq. (5.12) is a linear combination of four basis real timeseries, which in general are independent from each other. Note that this is not the case for the “standard” test Eq. (5.5), where only two independent timeseries are needed to describe the SNR timeseries. This is the direct consequence of the fact that in the “standard” case the signal model only depends on a single timeseries \hat{h}_{22}^R , while in the general case, the two timeseries h_+ , h_\times are needed to specify the signal model. We can gain more insight by redefining a different set of basis for $R_{\text{sym}}(t)$

$$\hat{h}_{++}^S(t) = \frac{1}{2} \left(\hat{h}_{++}(t) + \hat{h}_{\perp\perp}(t) \right) \quad (5.14)$$

$$\hat{h}_{++}^A(t) = \frac{1}{2} \left(\hat{h}_{++}(t) - \hat{h}_{\perp\perp}(t) \right) \quad (5.15)$$

$$\hat{h}_{+\perp}^S(t) = \frac{1}{2} \left(\hat{h}_{+\perp}(t) - \hat{h}_{\perp+}(t) \right) \quad (5.16)$$

$$\hat{h}_{+\perp}^A(t) = \frac{1}{2} \left(\hat{h}_{+\perp}(t) + \hat{h}_{\perp+}(t) \right) \quad (5.17)$$

and the predicted SNR timeseries takes a strikingly simple expression:

$$\begin{aligned} R_{\text{sym}}(t) = & z(0) \left(\hat{h}_{++}^S(t) + i\hat{h}_{+\perp}^S(t) \right) \\ & + z^*(0) \left(\hat{h}_{++}^A(t) + i\hat{h}_{+\perp}^A(t) \right). \end{aligned} \quad (5.18)$$

We can write Eq. (5.18) more compactly by introducing the two complex timeseries

$$\hat{h}^{S/A}(t) = \hat{h}_{++}^{S/A}(t) + i\hat{h}_{+\perp}^{S/A}(t), \quad (5.19)$$

so that

$$R_{\text{sym}}(t) = z(0) \hat{h}^S(t) + z^*(0) \hat{h}^A(t). \quad (5.20)$$

The expected SNR timeseries $R_{\text{sym}}(t)$ can be compared to the complex “symphony” SNR timeseries to yield a novel generally applicable signal consistency test ξ_{sym}^2 . We can compute the normalization of the ξ^2 statistic with a simple computation of the expected value of the residual timeseries $|z(t) - R_{\text{sym}}(t)|^2$ over different noise realizations:

$$E \left[\xi_{\text{sym}}^2(t) \right] = 2 - 2 \left| \hat{h}^S(t) \right|^2 + 2 \left| \hat{h}^A(t) \right|^2. \quad (5.21)$$

The expression generalizes Eq. (5.6) to the “symphony” search statistic. The details of the computation are reported in App. 5.A.

The first term in Eq. (5.18) has the same structure of the “standard” test Eq. (5.5). However, an additional term proportional to $z^*(0)$ enters the expression, thus breaking the convenient factorization between a complex trigger and a complex template-dependent timeseries. Of course, the expression agrees with the “standard” test for aligned-spin limit where $\hat{h}_+ = i\hat{h}_\times$. In that case, it is easy to show that $\hat{h}_\perp = \hat{h}_\times = \hat{h}_c$ and that

$$(\hat{h}_+ | \hat{h}_+)(t) = (\hat{h}_\times | \hat{h}_\times)(t) \quad (5.22)$$

$$(\hat{h}_+ | \hat{h}_\times)(t) = -(\hat{h}_\times | \hat{h}_+)(t). \quad (5.23)$$

For this reason, both $\hat{h}_{++}^A(t)$ and $\hat{h}_{+\perp}^A(t)$ are identically vanishing timeseries and Eq. (5.18) reduces to the “standard” case.

5.2.3 Approximating the New ξ^2 Test

Although the timeseries $R_{\text{sym}}(t)$ offers the best prediction for the “symphony” SNR, it cannot be expressed as a product of the trigger $z(0)$ and a template-dependent complex timeseries. As discussed earlier, such a convenient factorization is crucial for reducing the computational cost of the consistency test and for deploying the test with minimal changes to existing infrastructures. Therefore, we seek an expression

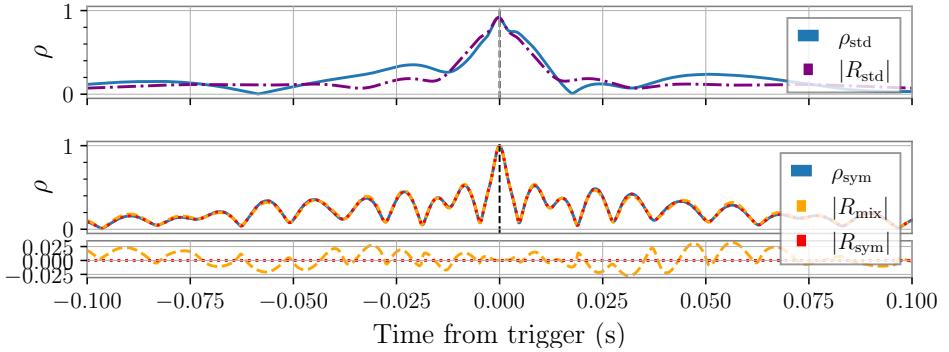


Figure 5.1: Predicted and measured absolute value of the SNR timeseries, $|R|$ and ρ respectively, for a precessing signal in zero noise and with unit magnitude, filtered with a perfectly matching template. We measure the SNR using both the “standard” SNR and the “symphony” SNR. The predicted SNR is computed with three different prescriptions, R_{std} Eq. (5.8) suitable for the “standard” SNR, and R_{sym} Eq. (5.18) and R_{mix} Eq. (5.24), both suitable for the “symphony” SNR. We also report in the bottom panel the difference between the expected and measured SNR timeseries. In the bottom panel we report the difference between both $|R_{\text{sym}}|$ or $|R_{\text{mix}}|$ and the “symphony” SNR. It is manifest that $|R_{\text{sym}}|$ and ρ_{sym} show perfect agreement between each other while $|R_{\text{std}}|$ and $|R_{\text{mix}}|$ do not accurately predict the relevant SNR timeseries. Also note that the peak of “standard” search statistic is lower than unity, meaning that performing matched filtering with templates Eq. (3.21) is not able to fully recover the SNR of a precessing signal. The signal is injected into Gaussian noise, sampled from the LIGO Livingston PSD [439] with a rate of 4096 Hz. The waveform is characterized by masses $m_1, m_2 = 28 \text{ M}_\odot, 3 \text{ M}_\odot$ and spins $\mathbf{s}_1 = (-0.8, 0.02, -0.5)$ and $\mathbf{s}_2 = 0$, observed with an inclination $\iota = 2.66$. It was generated starting from a frequency of 12 Hz with the approximant IMRPhenomXP [47].

$R_{\text{mix}}(t)$ for the predicted symphony SNR timeseries that retains this convenient factorization while providing satisfactory accuracy. While this expression is only an approximation to $R_{\text{sym}}(t)$, it may prove adequate in certain cases.

In particular, in the case of a precessing and HM binary system, the symmetries in Eqs. (5.22-5.23) are violated by only a “small amount”. More formally, we observe that the magnitude of $\hat{h}_{++}^A(t)$ and $\hat{h}_{+\perp}^A(t)$ is small in most of the practical cases and it makes sense to discard from Eq. (5.18) the term $\propto z^*(0)$. This leads to the an approximation $R_{\text{mix}}(t)$ of the predicted symphony timeseries:

$$R_{\text{mix}}(t) = z(0) \left(\hat{h}_{++}^S(t) + i \hat{h}_{+\perp}^S(t) \right) \quad (5.24)$$

Therefore, we can introduce the additional “mixed” signal consistency test ξ_{mix}^2 , which is obtained by comparing R_{mix} with the “symphony” SNR Eq. (5.10). The normal-

ization factor $E[\xi^2(t)]$ can be straightforwardly computed by setting $\hat{h}^A(t) = 0$ in Eq. (5.21):

$$E[\xi_{\text{sym}}^2(t)] = 2 - 2 |\hat{h}^S(t)|^2. \quad (5.25)$$

The test is equivalent to the “standard” ξ^2 Eq. (5.8) with the minimal replacement $\langle \hat{h}_p | \hat{h}_p \rangle(t) \rightarrow \hat{h}^S(t)$.

To obtain this expression we discarded from the predicted timeseries the complex quantity $z^*(0)\hat{h}^A(t)$. For this reason, the order of magnitude of $\hat{h}^A(t)$ is intimately connected with the error introduced by the “mixed” predicted timeseries, hence with the performance of ξ_{mix}^2 . In what follows we will use the magnitude of the peak ρ^A of the $\hat{h}^A(t)$ as a primary indicator of the goodness of the “mixed” consistency test

$$\rho^A = \max_t |\hat{h}^A(t)|. \quad (5.26)$$

As already noted, $\hat{h}^A(t)$ is identically zero for “standard” systems, hence ρ^A can be also used as a metric to quantify the amount of precession and/or HM of a template. Note that by definition $\max_t |\hat{h}^S(t)| = 1$, hence the quantity above automatically measures the ratio between \hat{h}^A and \hat{h}^S .

The choice of discarding the term $z^*(0)\hat{h}^A(t)$ is very natural but somehow arbitrary. Other choices for the residuals are possible, thus leading to a different expression for the predicted timeseries, hence different values of ξ^2 . A convenient alternative choice for an approximate signals consistency test consists in neglecting from $R_{\text{mix}}(t)$ all the terms $\mathcal{O}(\hat{h}_{+\times})$:

$$\begin{aligned} R_{\text{mix-bis}}(t) \simeq z(0) & \left[\frac{1}{2} \left(\hat{h}_{++}(t) + \hat{h}_{\times\times}(t) \right) \right. \\ & \left. + i \frac{1}{2} \left(\hat{h}_{+\times}(t) - \hat{h}_{\times+}(t) \right) \right] \end{aligned} \quad (5.27)$$

The expression has the merit of being more physically interpretable than Eq. (5.24), as it only depends on the physical polarizations, and it defines an additional test, labeled $\xi_{\text{mix-bis}}^2$. As we will see in the next section, the values of ξ^2 obtained with the latter expression do not significantly differ from the values of ξ_{mix}^2 obtained with Eq. (5.24), hence an experimenter interested in interpretability could freely use $\xi_{\text{mix-bis}}^2$ instead of ξ_{mix}^2 .

We close by noting that by assuming the *approximate* symmetries Eq. (5.22-5.23), $R_{\text{mix-bis}}(t)$ has the simple expression

$$R_{\text{mix-bis}}(t) \simeq z(0) \left(\hat{h}_{++}(t) + i \hat{h}_{+\times}(t) \right) \quad (5.28)$$

where we can straightforward recognize the “standard” test with the natural replacement $\hat{h}_p \rightarrow \hat{h}_+$ and $\hat{h}_c \rightarrow \hat{h}_\times$. This is the expression that one could have naively guessed without thorough computations.

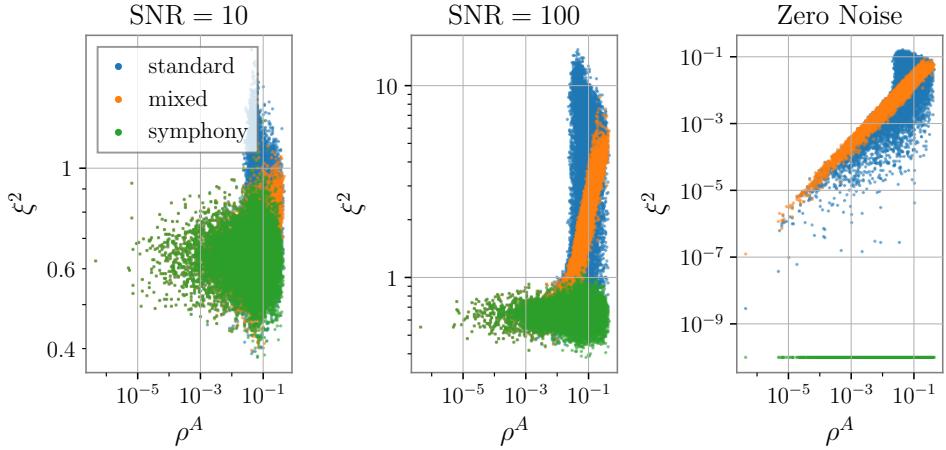


Figure 5.2: Values of ξ^2 Eq. (5.4) as a function of the absolute value of $(\hat{h}_+|\hat{h}_\times)$, which quantifies the precession and/or HM content of a signal. Each value is computed on random *precessing* BBHs, injected into Gaussian noise at a constant SNR. The left and center panels show SNRs of 10 and 100 respectively, while the right panel tackles the case of zero noise with injected signal normalized to 1. ξ^2 is computed using three different prescriptions. ξ_{std}^2 is obtained from $R_{\text{std}}(t)$ and $z_{\text{std}}(t)$ (label “standard”). ξ_{sym}^2 is computed using $R_{\text{sym}}(t)$ and $z_{\text{sym}}(t)$ (label “symphony”), while ξ_{mix}^2 uses $R_{\text{mix}}(t)$ and $z_{\text{sym}}(t)$. For this study we set ACL = 701.

In Fig. 5.1, we present an example demonstrating how accurately the predicted SNR timeseries aligns with the actual one, comparing the “standard,” “mixed,” and “symphony” cases. We compute the SNR timeseries obtained for precessing BBH signal with zero noise and filter the data with the same signal; we plot the measured SNR timeseries $\rho_{\text{std}}(t)$ and $\rho_{\text{sym}}(t)$ and the expected absolute values of the timeseries $R_{\text{std}}(t)$, $R_{\text{mix}}(t)$ and $R_{\text{sym}}(t)$. We note that $R_{\text{std}}(t)$ and $\rho_{\text{std}}(t)$ show poor agreement with each other, while $R_{\text{sym}}(t)$ perfectly models the complicated features of $\rho_{\text{sym}}(t)$. Moreover, $R_{\text{mix}}(t)$ provides a satisfactory approximation to $\rho_{\text{sym}}(t)$, accurate to a few percents.

5.3 Validity and Limitations of Different Signal Consistency Tests

In this section we study in depth the validity and the range of applicability of the various signals consistency tests discussed above, namely the new ξ_{sym}^2 , the “standard” ξ_{std}^2 and the approximated test ξ_{mix}^2 and its alternative expression $\xi_{\text{mix-bis}}^2$. After pre-

senting a study of the capabilities of the different tests, we study the performance of the “standard” and “mixed” tests as a function of the region of the parameter space. We also study the differences between the two alternative approximations of the “symphony” test ξ_{mix}^2 and $\xi_{\text{mix-bis}}^2$ and how the test depend on the choice of the integration window ACL. Finally, we compare the results obtained by performing the test in Gaussian noise with those obtained with real interferometer data: this study is crucial to test the robustness of the test in a “real life” scenario.

To carry out our analysis, we compute ξ_{std}^2 , ξ_{mix}^2 , $\xi_{\text{mix-bis}}^2$ and ξ_{sym}^2 for 15000 randomly sampled BBH signals, injected into Gaussian noise at different values of SNR. We uniformly sample the total mass $M \in [10, 50] \text{ M}_\odot$ and mass ratio $q = m_1/m_2 \in [1, 15]$, while reference phase and inclination, as well as the sky location, are drawn from a uniform distribution on the sphere. We also sample the starting frequency f_{min} in the range [5, 20] Hz. Conveniently, we can introduce the spin tilt angle θ , defined as

$$\theta_i = \arccos \frac{s_{iz}}{s_i}, \quad (5.29)$$

where s_i is the magnitude of the i -th spin. The tilt angle measures the mis-alignment of the each spin with the orbital momentum and thus it is crucial to control the amount of precession in a system, with a maximally precessing system having $\theta \simeq \pi/2$.

In our study, we explore two scenarios. In one case, we focus on precessing systems with the vectors’ end points of both spins homogeneously distributed inside the unit sphere. In the other case, we consider aligned-spin systems but include higher modes (HMs) in the waveform. For the latter experiment, we sample both z-components of the spins uniformly between $[-0.99, 0.99]$ and we consider the HM with $(\ell, |m|) = (2, 2), (2, 1), (3, 3), (3, 2), (4, 4)$. We utilize the frequency domain approximants IMRPhenomXP [47] and IMRPhenomXHM [422] for the two scenarios respectively. We employ the PSD computed over the first three month of the third observing run at the LIGO Livingston detector [439] and sample 100 s of Gaussian noise at a sample rate of 4096 Hz for each signal under study. To study the performance in real noise, we repeat the experiment by using segments of real publicly available real data [84], as we discussed with more details below.

5.3.1 How Does the Different Tests Compare to Each Other?

We compare here the performance of the different tests for different SNR summarizing the main results of the analysis described above. In figure Fig. 5.2, we present results pertaining precessing systems, while results in Fig. 5.3 refer to aligned-spin system with HM. In both figures for varying SNR, we plot the ξ^2 values as a function

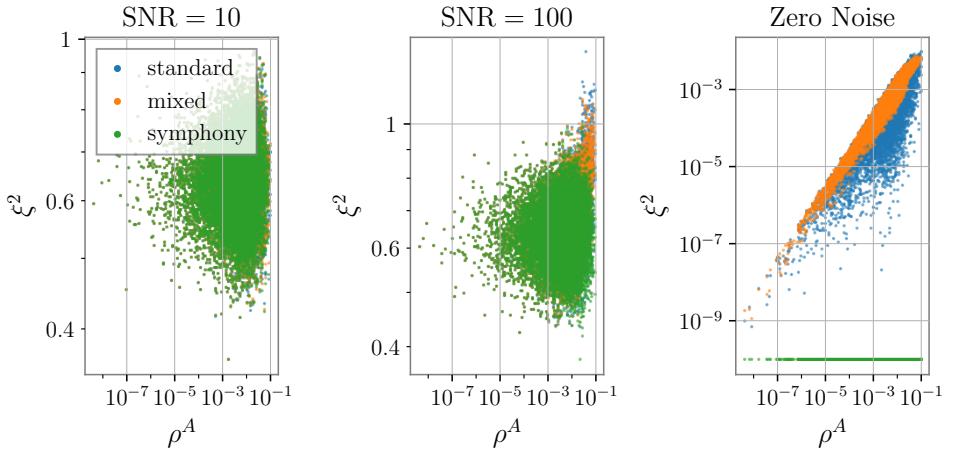


Figure 5.3: Values of ξ^2 Eq. (5.4) as a function of the absolute value of $(\hat{h}_+|\hat{h}_\times)$, which quantifies the precession and/or HM content of a signal. Each value is computed on random *aligned-spin* BBHs with HM content, injected into Gaussian noise at a constant SNR. The left and center panels show SNRs of 10 and 100 respectively, while the right panel tackles the case of zero noise with injected signal normalized to 1. ξ^2 is computed using three different prescriptions. ξ_{std}^2 is obtained from $R_{\text{std}}(t)$ and $z_{\text{std}}(t)$ (label “standard”). ξ_{sym}^2 is computed using $R_{\text{sym}}(t)$ and $z_{\text{sym}}(t)$ (label “symphony”), while ξ_{mix}^2 uses $R_{\text{mix}}(t)$ and $z_{\text{sym}}(t)$. For this study we set ACL = 701.

of the peak ρ^A of $\hat{h}^A(t)$, which as discussed above is an excellent measure of the content of precession and/or HM content in a template. Note that the ξ^2 values do not depend exclusively on ρ^A but also on the details of the residuals $\hat{h}^A(t)$. Nevertheless, ρ^A still remains a useful scalar quantity to quantify the lack of orthogonality of the two templates. In the “zero noise” case the injected signals are normalized to one: this arbitrary choice only affects the value of ξ^2 with an overall scaling but it does not alter the distribution of values. We compute the ξ^2 using a window (ACL) of 701 points centered around the injection time.

As long as zero noise is considered, we note that ξ_{std}^2 and ξ_{mix}^2 are both non-zero. This means that $R_{\text{std}}(t)$ and $R_{\text{mix}}(t)$ are not able to predict exactly the behaviour of the SNR timeseries. This is expected, since in the precessing/HM regime, they are both approximations to the true SNR. On the other hand, ξ_{sym}^2 is always zero (up to numerical noise), showing that the newly introduced ξ_{sym}^2 is the *optimal* test to search for generic BBH signals: this is also manifest in Fig. 5.1.

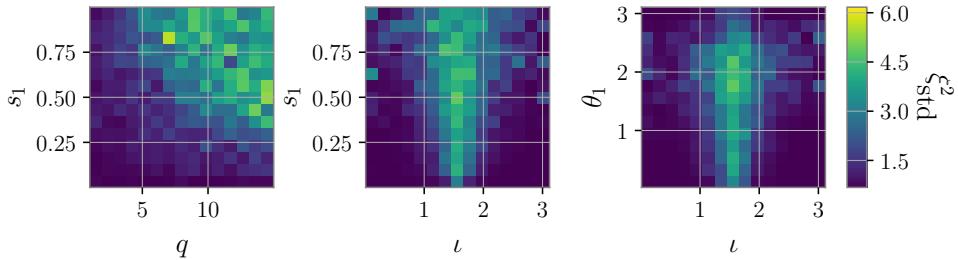


Figure 5.4: Performance of the “standard” test as a function of the parameter space. We color each bin according to the median value of ξ_{std}^2 and we consider the mass ratio q , the polar spin components s_1 , θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text.

Precessing templates By looking at the injected *precessing* signals in Fig. 5.2 in the presence of noise, we see that ξ_{sym}^2 is always superior to the “standard” test, with an improvement as large as two orders of magnitude in the $\text{SNR} = 100$ case. As long as ξ_{mix}^2 is considered, we observe that ξ_{mix}^2 and ξ_{sym}^2 show similar dispersions in the low SNR case. Therefore, in the presence of a substantial amount of noise, the accuracy improvement provided by ξ_{sym}^2 is negligible over the approximation given by ξ_{mix}^2 . The discrepancy between the two ξ^2 tests increases for $\text{SNR} = 100$; in that case, the noise level is lower and it must be of the same order of magnitude of the terms neglected to obtain ξ_{mix}^2 . We also note that the values ξ_{mix}^2 are very well correlated with ρ^A , thus confirming the usefulness of the latter to predict the failure of the “mixed” signal consistency test.

The fact that ξ_{mix}^2 degrades its performance at high SNR should not be of concern, as the signal consistency test is less crucial for the high SNR region. Indeed, due to the rarity of very loud signals, it is feasible to perform targeted follow up and *ad hoc* studies, hence assessing the significance of a trigger with other strategies. For this reason, we conclude that ξ_{mix}^2 is likely to perform close to optimality for the vast majority of the practical applications and we recommend its implementation in any pipeline aiming to search for precessing signals.

In closing, we note that in the ξ_{std}^2 computation, we could have used Eq. (3.26) instead of Eq. (5.7) to filter the data and, similarly, Eq. (5.5) rather than Eq. (5.8) to compute the autocorrelation. Even though in the “standard” case, the two expressions agree, they do not agree when precessing and/or HM template are considered. In that case, a straightforward generalization to precession $\hat{h}_p \rightarrow \hat{h}_+$ and $\hat{h}_c \rightarrow \hat{h}_\times$ would lead ξ_{std}^2 to also take into account both polarizations, with potential improvements. However, as the GstLAL pipeline implements the test using Eq. (5.8), we made the choice to describe the current situation.

HM templates The picture outline above changes when aligned-spin HM templates are considered in Fig. 5.3. In this case, the performance of ξ_{std}^2 , ξ_{mix}^2 and ξ_{sym}^2 are very comparable in the low SNR case. In the high SNR case, ξ_{sym}^2 retains a slightly better performance but the use of ξ_{mix}^2 does not bring any additional improvement over the “standard” test ξ_{std}^2 . Therefore, if only HMs are considered, our results indicate that ξ_{std}^2 already delivers close to optimal results, suggesting that no updates of the “standard” signal consistency test are required to tackle only the aligned-spin HM case.

5.3.2 When Does the “Standard” Test Fail?

To study the limitation of the “standard” ξ^2 , in Fig. 5.4 we report the values of ξ_{std}^2 as a function of the template parameters. We only focus on the *precessing case*, because in the HM case ξ_{std}^2 Fig. 5.3 shows a good performance across the whole parameter space. Our results suggest that the performance of the test decreases for large values of the mass ratio, large values of spin and for large spin misalignment. Moreover, precession is more visible for systems observed with a close to edge-on inclination, i.e. $\iota \simeq \pi/2$. These findings align with the literature on the detectability of precession in BBH [73, 74, 182, 406], which demonstrates that precession is more detectable for asymmetric, heavily spinning edge-on systems. The GW signals emitted by such heavily precessing acquire a more complicated structure, which translates into the lack of symmetry between the two polarizations, which fail to satisfy Eq. (5.1). As the “standard” test relies on such symmetry and uses only the plus polarization to predict the SNR timeseries, a large violation of Eq. (5.1) also leads straightforwardly to large values of ξ_{std}^2 .

5.3.3 How Does the “Mixed” Test Perform?

As we turn our attention to the performance of the approximate “mixed” signal consistency test, we are interested in (i) identifying the regions of the parameter space where this test provides an advantage over the “standard” test and (ii) evaluating the extent of performance loss compared to the optimal “symphony” test. We limit our study to precessing signals, since for aligned-spin HM systems the three tests show very similar performance.

To answer the first question, we report in Fig. 5.5 the ratio $\xi_{\text{mix}}^2/\xi_{\text{std}}^2$ between the “mixed” and the “standard” test, evaluated at SNR = 100 as a function of the template parameters. Unsurprisingly the “mixed” test outperforms the “standard” test for systems that show a strong amount of precession: in these regions of the parameter space the “standard” test shows poor performance while the “mixed” test is able

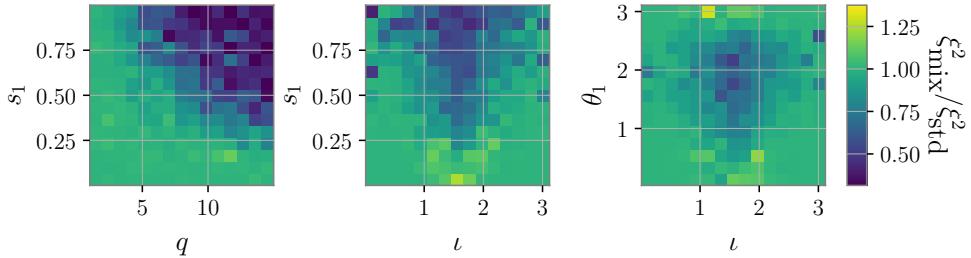


Figure 5.5: Comparison between the “mixed” and the “standard” signal consistency tests as a function of the parameter space. We color each bin according to the median value of $\xi_{\text{mix}}^2 / \xi_{\text{std}}^2$ and we consider the mass ratio q , the polar spin components s_1, θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text.

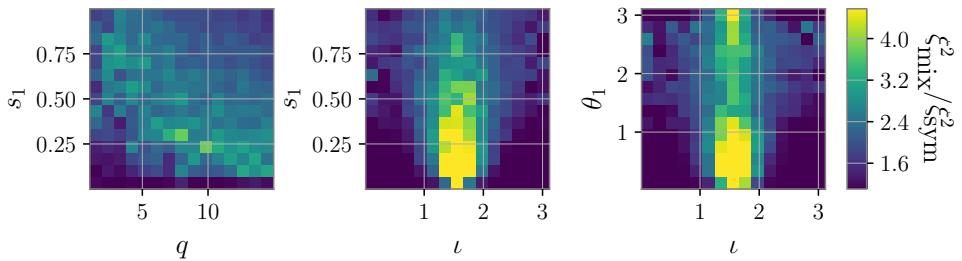


Figure 5.6: Comparison between the “mixed” and the “standard” signal consistency tests as a function of the parameter space. We color each bin according to the median value of $\xi_{\text{mix}}^2 / \xi_{\text{sym}}^2$ and we consider the mass ratio q , the polar spin components s_1, θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text.

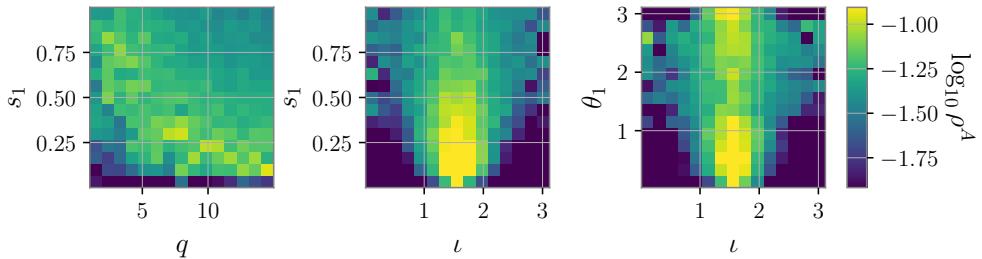


Figure 5.7: Values of ρ^A as a function of the parameter space. We color each bin according to the median value of ρ^A and we consider the mass ratio q , the polar spin components s_1, θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text. ρ^A is intended as a measure of the goodness of the “mixed” signal consistency test and as such, it correlate very well with the values of ξ_{mix}^2 (see also Fig. 5.6).

to better predict the behaviour of the SNR timeseries.

To understand the performance loss of the “mixed” test as compared to the optimal “symphony” test, we turn our attention to Fig. 5.6, where we report the ratio $\xi_{\text{mix}}^2 / \xi_{\text{sym}}^2$, also evaluated at $\text{SNR} = 100$, as a function of the template parameters. By comparing the performance of the ξ_{mix}^2 with ξ_{sym}^2 it striking to note that the “mixed” test has poor performance for systems with $\iota \simeq \frac{\pi}{2}$. This is somehow expected as it is well known that the inclination increases the precession induced amplitude and phase modulation on the waveform. However, it is surprising that the worst performance occurs mostly for systems with low precession content, i.e. with $s_1 \lesssim 0.25$ and for $\theta_1 \lesssim 1$. This is puzzling because the precession content of such systems is expected to be small, due to low spins and mild mis-alignment. We make the hypothesis that the effect can be explained by studying the details of the *spin twist* procedure [171, 172, 180] employed to model the precessing effects in the approximant in use. According to this procedure, the waveform is generated with a time dependent rotation of the waveform emitted by a corresponding aligned-spin system. The spin twist might introduce some numerical noise in the polarizations, even in the aligned-spin limit. For ι far from $\frac{\pi}{2}$, the numerical noise is expected to be much smaller than the waveform itself, however as $h_x \rightarrow 0$ for $\iota \rightarrow \frac{\pi}{2}$ the noise might become dominant and affect the computation of $\hat{h}^S(t)$. A numerical study reported in Fig. 5.8 shows that even in the *aligned-spin* case the precessing approximant `IMRPhenomXP` returns a non-zero value for ρ^A , while the aligned-spin approximant `IMRPhenomXAS` [158] returns the zero value consistent with the theory. For this reason, as `IMRPhenomXP` is built upon `IMRPhenomXAS`, we may conclude that the spin twist procedure introduces spurious noise into the polarizations. The numerical noise becomes visible in the timeseries $\hat{h}^A(t)$ only when the cross polarization tends to zero, which happens for close to aligned-spin systems but not for heavily precessing systems: that explains why the issue is seen only for small values of s_1 . Further investigations should corroborate this hypothesis.

These limitations for $\iota \simeq \frac{\pi}{2}$, most likely due to the approximant employed, should not be of concern. First of all, they affect a region of the parameter space well covered by current searches and for which the “standard” test is mostly suitable. Secondly, future precessing approximants might solve the pathological behaviour observed in Fig. 5.8, thus potentially improving the performance of the “mixed” test. Finally, we note that ξ_{mix}^2 differs from ξ_{sym}^2 by at most a factor of ~ 2 . It is left to the developer of a search to quantify the consequent loss in sensitivity and to consider whether this is an acceptable loss. Here we limit ourself to stress that ξ_{mix}^2 is more suitable to deal with the heavily precessing regime and provides a substantial improvement over the “standard” test. Moreover, we stress again that the results discussed refers to signals

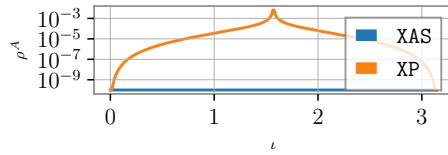


Figure 5.8: Numerical stability of the approximant `IMRPhenomXP` in the aligned-spin limit as a function of the inclination angle ι . We compute the values of ρ^A and for a BBH with total mass $M = 10 M_\odot$ and mass ratio $q = 8$ and with spins $s_{1z}, s_{2z} = -0.4, 0.6$, seen at different inclination angles ι . We repeat the experiment with both the precessing approximant `IMRPhenomXP` and the aligned-spin approximant `IMRPhenomXAS`. While for the aligned-spin approximant ρ^A are both zero up to numerical precision, for the precessing approximant the two quantities become non zero as the inclination gets close to $\iota \simeq \pi/2$. In this case, since $h_x \rightarrow 0$, the numerical noise introduced by the spin twist procedure affects the normalized cross polarization \hat{h}_x .

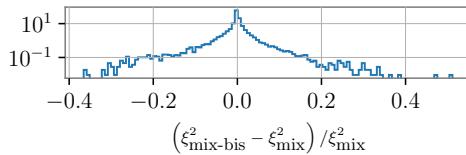


Figure 5.9: Discrepancies between the “mixed” signal consistency test ξ_{mix}^2 Eq. (5.24) and a possible alternative definition Eq. (5.27) $\xi_{\text{mix-bis}}^2$. For the 15000 BBH signals introduced in the text, we report a histogram with the fractional difference $\frac{\xi_{\text{mix-bis}}^2 - \xi_{\text{mix}}^2}{\xi_{\text{mix}}^2}$ between the two consistency tests. The discrepancies between the two versions are negligible, showing that Eq. (5.24) and Eq. (5.27) are mostly equivalent.

with a high SNR of 100, for which a *ad hoc* follow up strategy can be implemented, thus making the signal consistency test less decisive. For signals observed with a more realistic lower SNR, the impact on the ξ_{mix}^2 values is much smaller, as shown in Fig. 5.2.

In Fig. 5.7 we report the peak value ρ^A of the residual timeseries $\hat{h}^A(t)$ as a function of the template parameters. We observe that the values of ρ^A are very well correlated with the values of ξ_{mix}^2 reported in Fig. 5.6. This confirms the utility of ρ^A as a measure of the performance of the ξ_{mix}^2 .

5.3.4 Which “Mixed” Test Should We Use?

As a final analysis, we compare the performance of the two approximate tests introduced in Sec. 5.2.3, ξ_{mix}^2 and $\xi_{\text{mix-bis}}^2$. While ξ_{mix}^2 has a more straightforward definition,

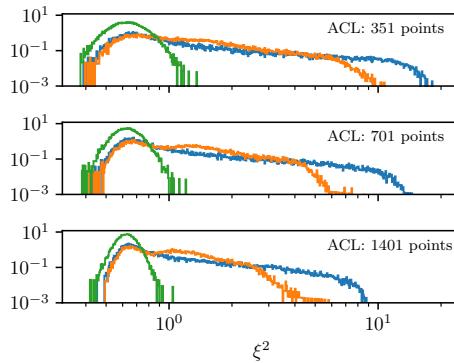


Figure 5.10: Values of ξ^2 for different window lengths (ACL) used for the integral in Eq. (5.4). For each ACL, we report the values of ξ_{std}^2 (blue), ξ_{mix}^2 (orange) and ξ_{sym}^2 (green), following the color code introduced in Figs. 5.2 and 5.3. The data refers to 15000 *precessing* signals injected at $\text{SNR} = 100$.

$\xi_{\text{mix-bis}}^2$ has a simpler expression and it is only written in terms of the two polarizations. Depending on the pipeline details, an user may decide to implement either versions. For this reason, it is important to check that they give consistent results. This is done in Fig. 5.9 where we report an histogram with the relative error of $\xi_{\text{mix-bis}}^2$ with respect to ξ_{mix}^2 . From the histogram we learn that in 90% of the cases $\xi_{\text{mix-bis}}^2$ has a discrepancy of less than 6% from ξ_{mix}^2 . Our results show that the two expressions for the predicted SNR timeseries, Eq. (5.24) and Eq. (5.27), are mostly equivalent and they can both successfully employed in a full search, with the caveats about the validity of the approximation discussed above.

5.3.5 How to Choose the Autocorrelation Length for the Test?

In Fig. 5.10, we study how the three ξ^2 tests considered depend on the choice of the integration window ACL. For brevity, we only consider the ξ^2 values for *precessing* injections with $\text{SNR} = 100$ and we choose three different values of $\text{ACL} = 351, 701, 1401$. These values are routinely employed by different analysis carried by the GstLAL pipeline and were chosen after careful numerical simulations [77].

Our results show that the ξ^2 values are slightly improved for a longer integration window, i.e. for a larger ACL. However, the smallness of the differences suggests that the choice of ACL is not crucial and that overall the test is robust against different ACL values. We then recommend to choose $\text{ACL} = 701$, which seems a satisfying trade-off between test performance and computational cost.

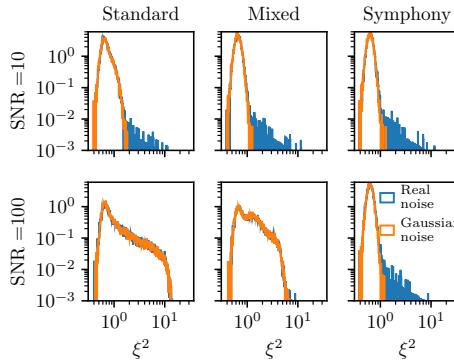


Figure 5.11: Performance in real detector’s noise of the three signal consistency tests discussed in the paper. We randomly select a number of precessing templates as in Fig. 5.2 and we inject a corresponding signal into both Gaussian noise and real detector’s noise. The detector’s noise was recorded by the LIGO-Hanford observatory between GPS times 1245708288 s and 1246756864 s. In each histogram, we report the values of ξ^2 for the Gaussian noise (orange) and Real noise (blue) case. Each panel refers to a different type of signal consistency test and to different values of the injected SNR.

5.3.6 How Does Real Noise Affect the Test Performance?

We conclude our analysis by studying the performance of the various signal consistency tests in *real detector’s noise*. This study allows us to study a realistic scenario, where non-Gaussian artifacts may negatively impact our ability to predict the SNR timeseries. To do so, we use the publicly available data [84] taken during the third Observing Run (O3) by the LIGO-Hanford observatory between GPS times 1245708288 s and 1246756864 s. We use the data to produce *whitened* segments of 100 s, where we inject the test signals introduced above. We then compute the SNR timeseries, both for the “standard” and the “symphony” case, and we compute the ξ^2 for each injection. We present our results in Fig. 5.11, where for each type of signal consistency test we report the distribution of ξ^2 values computed in both the real and Gaussian noise cases. As above, we repeat the experiment for signals injected at an SNR of 10 and 100.

From our results it is manifest that the bulk of the ξ^2 distributions computed for both real and Gaussian noise are consistent with each other: this confirms the robustness of the test even in the real noise scenario. Moreover, we note that our novel “symphony” produces systematically lower values of ξ^2 as compared to both the “mixed” and “standard” tests, further confirming its optimal performance in the high SNR case.

In some occasions the real noise produces triggers with large values of ξ^2 when compared with the Gaussian noise case: in these situations the non-Gaussianities of the detector's output introduce artifacts in the SNR timeseries, which are not taken into account by our prediction. This produces the large values of ξ^2 observed in Fig. 5.11. This behaviour is not only expected but also *desired*: indeed, the ξ^2 test is specifically designed to discriminate between astrophysical triggers and triggers from noise artifacts. Consequently, large ξ^2 values in the presence of noise artifacts indicate that the test is correctly performing its intended function.

For triggers at $\text{SNR} = 100$, the “standard” and “mixed” signal consistency tests agree when computed for Gaussian and real noise: in this case, the test in Gaussian noise already produces triggers with large ξ^2 values, which are indistinguishable from the effect of non-Gaussian artifacts.

5.4 Final Remarks

We introduce a novel ξ^2 signal consistency test tailored for the matched-filtering searches of gravitational waves emitted by precessing and/or higher-mode binary systems. The test measures the discrepancy between the predicted and measured SNR timeseries, as output by any matched-filtering pipeline which filters the data with a large set of templates. While the traditional test ξ_{std}^2 is only valid for the case of aligned-spin binaries where HM are not considered, our new ξ_{sym}^2 is built upon the “symphony” search statistic [328] and it does not make any assumption about the nature of the signal to detect. Our new ξ_{sym}^2 relies on the expression Eq. (5.18) for the predicted matched-filtering output, which is derived here for the first time.

Thanks to the symmetry of the aligned-spin systems without HMs, the “standard” consistency test is conveniently factorized as the multiplication between a complex trigger and a complex template-dependent timeseries. As our newly introduced test breaks such simple factorization, it requires twice the computational cost to be performed. To alleviate such cost, we also proposed an approximation ξ_{mix}^2 to our new ξ_{sym}^2 , which, by respecting the simple factorization for the predicted timeseries, requires minimal changes to existing code platforms and is obtained without extra computation with respect to the “standard” test.

We investigated the validity of the two newly introduced tests ξ_{sym}^2 and ξ_{mix}^2 , by performing an extensive study on signals injected in both Gaussian noise, reaching to four main conclusions:

- The newly introduced test ξ_{sym}^2 has optimal performance for precessing and/or HM systems, being able to optimally predict the SNR timeseries in the zero noise case

- The traditional test ξ_{std}^2 show poor performance for heavily asymmetric systems with mis-aligned spins
- The approximate test ξ_{mix}^2 is very suitable for low SNR signals, while it displays some loss of performance for systems with high SNR
- The traditional test ξ_{std}^2 is very suitable for aligned-spin systems where HM are considered, showing similar performances to the optimal ξ_{sym}^2

These same conclusions are obtained by studying the test performance on real detector's noise, although with some outliers with high ξ^2 values in correspondence of loud non-Gaussian transient burst of noise.

The newly introduced test, ξ_{sym}^2 , can be implemented in any matched-filtering pipeline, enhancing the search for both precessing and aligned-spin signals with HM content. Additionally, the validation studies presented here will help the community better understand the limitations of the traditional ξ_{std}^2 test, benefiting any matched-filtering pipeline aimed at detecting binary black hole signals with strong precession and/or higher-mode content.

5.A Expected Value of the “Symphony” Signal Consistency Test in Gaussian Noise

To compute the expected value of the “symphony” signal consistency test in Gaussian noise we closely follow the computation presented in [77, App. A]. We begin by noting that since the two templates \hat{h}_+ , \hat{h}_\perp and the data are real timeseries, the matched-filtering output $\langle d|h_{+/\times}\rangle(t)$ is also a real timeseries: $\langle d|h_{+/\perp}\rangle(t) = (d|h_{+/\perp})(t)$. Thus, the complex SNR timeseries $z_{\text{sym}}(t) = (d|\hat{h}_+)(t) + i(d|\hat{h}_\times)(t)$ can be re-written as:

$$z_{\text{sym}}(t) = \langle d|\hat{h}_+ + i\hat{h}_\perp\rangle(t). \quad (5.30)$$

We can now compute the expect value of the residual timeseries $E\left[\xi_{\text{sym}}^2(t)\right]$ by

taking the ensemble average of $|z(t) - R_{\text{sym}}(t)|^2$:

$$\begin{aligned}
 \mathbb{E} [\xi_{\text{sym}}^2(t)] &= \mathbb{E} \left[\left| z(t) - z(0) \hat{h}^S(t) - z^*(0) \hat{h}^A(t) \right|^2 \right] \\
 &= -2\Re \left\{ \mathbb{E} [z^*(t)z(0)] \hat{h}^S(t) \right\} \\
 &\quad - 2\Re \left\{ \mathbb{E} [z(t)z(0)] \hat{h}^{A*}(t) \right\} \\
 &\quad + 2\Re \left\{ \mathbb{E} [z(0)z(0)] \hat{h}^{*A}(t) \hat{h}^S(t) \right\} \\
 &\quad + \mathbb{E} [|z(0)|^2] \left| \hat{h}^S(t) \right|^2 + \mathbb{E} [|z(0)|^2] \left| \hat{h}^A(t) \right|^2 \\
 &\quad + \mathbb{E} [|z(t)|^2]
 \end{aligned} \tag{5.31}$$

To move forward, we need to consider some properties of Gaussian noise in frequency domain, introduced in Sec. 2.2.2:

$$\mathbb{E} [\tilde{n}^*(f') \tilde{n}(f)] = S(f) \delta(f - f') \tag{5.32}$$

$$\mathbb{E} [\tilde{n}(f') \tilde{n}(f)] = 0. \tag{5.33}$$

Using these two properties it is easy to show that

$$\mathbb{E} [|z(t)|^2] = \mathbb{E} [|z(0)|^2] = 2 \tag{5.34}$$

$$\mathbb{E} [z(0)z(0)] = \mathbb{E} [z(t)z(0)] = 0 \tag{5.35}$$

by computing the relevant integrals. Moreover, we can easily compute $\mathbb{E} [z^*(0)z(t)]$ using the definition of scalar product Eq. (3.36):

$$\begin{aligned}
 \mathbb{E} [z^*(0)z(t)] &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} df df' \frac{\tilde{n}^*(f') \tilde{n}(f)}{S(f') S(f)} \times \\
 &\quad \times (\hat{h}_+^* - i\hat{h}_\perp^*)(\hat{h}_+ + i\hat{h}_\perp) e^{-i2\pi f t} \\
 &= \int_{-\infty}^{\infty} df \frac{\hat{h}_+^* \hat{h}_+ + \hat{h}_\perp^* \hat{h}_\perp + i(\hat{h}_+^* \hat{h}_\perp - \hat{h}_\perp^* \hat{h}_+) }{S_n(f)} e^{-i2\pi f t} \\
 &= 2\hat{h}^S(t).
 \end{aligned} \tag{5.36}$$

Putting everything together, we obtain a simple expression for $\mathbb{E} [\xi_{\text{sym}}^2(t)]$:

$$\mathbb{E} [\xi_{\text{sym}}^2(t)] = 2 - 2 \left| \hat{h}^S(t) \right|^2 + 2 \left| \hat{h}^A(t) \right|^2 \tag{5.37}$$

which can be used to normalize the “symphony” ξ^2 test. Note that in the “standard” limit, we have $\hat{h}^A(t) \rightarrow 0$ and $\hat{h}^S(t) \rightarrow \langle \hat{h}_p | \hat{h}_p \rangle(t)$, hence the two expected values $\mathbb{E} [\xi^2(t)]$ agree.

CHAPTER 6

Searching for Gravitational-wave Signals from Precessing Binaries

6.1 Introduction

As already discussed, the majority of matched-filtering pipelines heavily relies on the simplicity of aligned-spin systems to obtain a convenient expression for the matched-filter and to reduce the number of templates required to cover the space of interest. In this Chapter, we move on from a search for aligned-spin systems and we introduce a search method suitable for precessing systems. The search is particularly useful to target asymmetric binaries with significantly misaligned spins, where a substantial sensitivity reduction is expected [73–75]. The method presented lay the foundation for a more in-depth study of precessing binaries, with implications for our understanding of binary formation channels and gravitational wave cosmology.

Several approaches have been developed to tackle the challenges of precessing searches. Seminal works [42, 334, 440] focused on modeling precession with a set of phenomenological parameters, which were subsequently used to perform a matched-filtering search [441]. In this approach, the templates do not correspond to real physical signals but they are designed to match adequately physical precessing signals. More recently and in a similar fashion, the authors of [406] leveraged the so-called two-harmonic approximation of a precessing signal [185] to achieve a substantial sensitivity increase at a moderate computational cost when searching for precessing signals in the NSBH parameter space.

A different approach consists in directly using the physical waveform to filter the data [66, 327, 328]. This corresponds to a straightforward extension of the traditional matched-filter pipeline and has the benefits of using the real physical templates and, as such, it does not rely on any approximation of the waveforms, being applicable to virtually any type of binary system (including non-GR waveforms). Employing this method typically requires a large number of templates and, as such, it poses the challenge of effectively generating large, high-dimensional template banks at a feasible computational cost. In Ch. 4 we addressed this challenge and developed a novel template placement method, specifically designed to generate large template banks, with minimal assumptions on the nature of the GW signal. Hence, by cutting the bank generation time, the exploration of the heavily precessing parameter space was made feasible.

Similar to the direct filtering method of [328], here we design a search method for precessing signals using the features of the GstLAL search pipeline. Our method was first presented in [3]. Besides using the template placement algorithm of Ch. 4, our method relies on the results of Ch. 5 for an effective signal consistency test. We demonstrate that our method results in an improvement in the pipeline’s sensitivity to systems with large spin-orbit misalignment, reaching up to 100% increase in the search sensitive volume, depending on the region of the parameter space. Furthermore, our findings suggest that the primary cause of degraded performance in an aligned-spin search targeting precessing signals is not a poor SNR recovery but rather the failure of the ξ^2 signal consistency test. We stress that with the new upgrades, the GstLAL pipeline is ready to handle not only a search for precessing signals but also a search for generic BBH signals, such as those generated by eccentric binaries [68, 399] or those with imprints from higher-order modes (HMs) [105, 373, 407].

In Sec. 6.2 we describe our search method and the modifications to the GstLAL pipeline needed to achieve our goal, while in Sec. 6.3, we generate and validate two template banks, targeting heavily precessing BBHs. We use these banks to conduct two searches using real data from the LIGO and Virgo interferometers during the third observing run (O3) [83, 84]. In Sec. 6.4, we present the results from these searches, validate their performance and report the sensitivity improvement in comparison to the aligned-spin version of the pipeline. In Sec. 6.4, we describe future work which will further improve the pipeline’s sensitivity to precessing binary systems. Final remarks are gathered in Sec. 6.6.

6.2 Search Method

As discussed in Sec. 1.3.3, the rotation of the orbital plane induced by precession [39] breaks the equatorial symmetry of the system. This translates into the breaking of the symmetry between the $+$ and the \times frequency-domain polarizations \tilde{h}_+ , \tilde{h}_\times of the GW generated by an aligned-spin systems which is typically expressed as

$$\tilde{h}_+ \propto i\tilde{h}_\times. \quad (6.1)$$

As most of the aligned-spin search pipelines, including GstLAL, use this symmetry to simplify the search problem [63, 77, 355, 360], the development of a precessing search pipeline must relax the assumption of Eq. (6.1) and address the general case. This entails the introduction of the general search statistic discussed in Sec. 3.2.2 and of the modified signal consistency test discussed in Ch. 5. In what follows we summarize how such upgrades are implemented in the GstLAL pipeline.

General search statistic In this work, we compute the SNR timeseries for a precessing template as in Eq. (3.33):

$$z_{\text{prec/HM}}(t) = (s|\hat{h}_+)(t) + i(s|\hat{h}_\perp)(t) \quad (6.2)$$

where we introduced the “orthogonalized” template¹:

$$\hat{h}_\perp(t) = \frac{1}{\sqrt{1 - \hat{h}_{+^\times}^2}} \left(\hat{h}_\times(t) - \hat{h}_{+^\times} \hat{h}_+(t) \right). \quad (6.3)$$

This expression is particularly convenient as it enables the straightforward implementation of the search statistic in the GstLAL pipeline by simply modifying the filter waveforms.

It is worth noting that an alternative search statistic, suitable only for precessing signals, is introduced in [327]. While the alternative statistic has the obvious advantage of also maximizing over the reference phase ϕ , thus reducing the dimensionality of the template banks. it comes at the price of a more complex functional form which translates into an increased computational cost of the search. Moreover, the statistic [327] can be applied in a less general scenario as it does not apply to signals with HM content. Indeed by implementing the statistic Eq. (6.2), our pipeline is ready to search for signals with imprints from HMs, thus reproducing the successful search presented in [373]. See however [409, 442, 443] for a promising alternative method to search for systems with HMs, which does not rely on Eq. (6.2).

¹Note that, trivially, $(\hat{h}_\perp|\hat{h}_\perp) = 1$.

General ξ^2 test As discussed in Ch. 5, in the aligned-spin case the predicted SNR timeseries $R_{\text{std}}(t)$ for a given template has a simple expression

$$R_{\text{std}}(t) = (\hat{h}_+ | \hat{h}_+)(t) + i(\hat{h}_+ | \hat{h}_+^{\pi/2})(t). \quad (6.4)$$

In Sec. 5.2.2 we showed that in the precessing case the true template response is given by Eq. (5.18) and it depends on both the trigger $z(0)$ and on its complex conjugate $z^*(0)$, thus losing the simple factorization of $R(t)$ as a product between the trigger $z(0)$ and a complex template-dependent timeseries. This translates into an increased computational cost, and implementing the exact expression would require heavy changes to the pipeline.

Luckily, in Sec. 5.3 we showed that for $\text{SNR} \lesssim 100$ an *approximate* template response is adequate for the purpose of the ξ^2 test, while at the same time keeping the simple simple factorization of the standard template response. Therefore, since a signal has typically a much lower SNR, the use of the approximation is justified and we can use the “mixed” signal consistency test we introduced earlier in Sec. 5.2.3. Thus, the predicted SNR timeseries suitable for our precessing pipeline is given by:

$$\begin{aligned} R_{\text{prec}}(t) = & z(0) \frac{1}{2} \left[(\hat{h}_+ | \hat{h}_+)(t) + (\hat{h}_\times | \hat{h}_\times)(t) \right] \\ & + i \frac{1}{2} \left[(\hat{h}_+ | \hat{h}_\times)(t) - (\hat{h}_\times | \hat{h}_+)(t) \right]. \end{aligned} \quad (6.5)$$

We note that we choose to implement in the pipeline the expression Eq. (5.27), as opposed to the one in Eq. (5.24). While the two expression are mostly equivalent and provide very similar results, our chosen expression has the advantage of being more physically interpretable. Based on our study in Sec. 5.3.4, we do not expect this choice to alter the search sensitivity.

Ranking statistic For each trigger happening in any of the instruments $\{H1, L1, V1\}$, the GstLAL pipeline records the trigger time t , the phase Φ_T of the trigger, the SNR ρ , the ξ^2 value and the instantaneous sensitivity of each detector, measured in terms of horizon distance D . Triggers happening on different instruments within a similar time window and with the same template are grouped together, forming a coincident trigger or *coincidence*.

As discussed in Sec. 3.4.3, to discriminate between coincidences arising from random fluctuations and coincidences of astrophysical origin, it is crucial to develop an accurate ranking statistic. Within the GstLAL framework, this is accomplished by using the ranking statistic Eq. (3.54) for a coincidence $\{\vec{D}, \vec{\rho}, \vec{\xi^2}, \vec{t}, \vec{\Phi}_T\}$:

$$\mathcal{L} = \frac{p(\{\vec{D}, \vec{\rho}, \vec{\xi^2}, \vec{t}, \vec{\Phi}_T\} | \text{signal})}{p(\{\vec{D}, \vec{\rho}, \vec{\xi^2}, \vec{t}, \vec{\Phi}_T\} | \text{noise})}. \quad (6.6)$$

When it comes to precessing searches, the trigger time t , SNR, ξ^2 and D keep the same physical meaning they have in the standard case and, for this reason, the terms in the LR involving only such quantities may not require any modification for the precessing case. On the other hand, due to a different signal model, the recorded phase Φ_T loses its straightforward physical interpretation it has in the aligned-spin case (see also Eq. (3.23)). Therefore, the factors of \mathcal{L} depending on ϕ may require to be adapted for a precessing search. The phase dependent terms measure the coherence of the triggers among multiple detectors and in the GstLAL pipeline they correspond to $p(\vec{p}, \vec{t}, \vec{\Phi}_T | \vec{D}, \text{signal})$ and $p(\vec{t}, \vec{\Phi}_T | \vec{D}, \text{noise})$.

Since we have no reason to believe that the ranking statistic, aside from the “coherence” term, is unsuitable for the precessing case, in this study we maintain without modification the ranking statistic utilized by the pipeline in its latest version [82]. However, the use of “coherence” terms designed for aligned-spin searches in the precessing case could potentially impact the search sensitivity negatively. Future work may investigate this impact on sensitivity and could develop an improved ranking statistic by computing and implementing an expression for the “coherence” term that is suitable for the precessing case.

6.3 Template Banks

To demonstrate the effectiveness of the updated GstLAL pipeline, we conduct two runs employing large template banks constructed for precessing signals. One is tailored for mildly asymmetric systems (referred to as the “Low q ” bank), while the other focuses on more asymmetric systems (the “High q ” bank). For each precessing search, we run an aligned-spin search targeting the same mass range. This comparative approach enables an assessment of the enhanced sensitivity resulting from the use of precessing templates and the modifications in the filtering scheme described in the previous section. The next section will present the results of said comparison. Here, we delve into the template banks’ generation, providing insight into our choices for parameter space and validating their ability to cover the target regions. A summary of the template banks generated is presented in Tab. 6.1.

As it is standard in the literature, the size and performance of the bank are controlled by the *target minimal match* MM [336]: as mentioned in Sec. 3.3, a good bank should have only a small fraction $\lesssim 10\%$ of random signals with a match lower than MM with the nearest template of the bank. The actual fraction of signals with a match below MM depends on the details of the template bank generation as well as on the parameter space of choice.

To generate the *aligned-spin* template banks, we use the state-of-the-art *stochastic*

	Mass range	Parameter space	AS	Size P
Low q	$m_1, m_2 \in [8, 70]M_\odot$ $q \in [1, 6]$	$s_{1z} \in [-0.99, 0.99]$ $s_{2z} \in [-0.99, 0.99]$ --- $s_1 \in [0, 0.9]$ $\theta_1 \in [-\pi, \pi]$ $s_{2z} \in [-0.99, 0.99]$ $\iota \in [0, \pi]$	8425	1605625
High q	$m_1 \in [15, 70]M_\odot$ $m_2 \in [3, 10]M_\odot$ $q \in [5, 12]$	$s_{1z} \in [-0.99, 0.99]$ $s_{2z} \in [-0.99, 0.99]$ --- $s_1 \in [0.5, 0.9]$ $\theta_1 \in [-\pi, \pi]$ $s_{2z} \in [-0.99, 0.99]$ $\iota \in [0, \pi]$	27016	2287083

Table 6.1: We summarize here the most important features of the four template banks considered in this work. We select two regions of the parameter space, labeled “Low q ” and “High q ” respectively. For each region, we generate two template banks, one only gathering aligned-spin signals (AS) and another one including precessing signals (P). The aligned-spin banks covers the variable M, q, χ_{eff} , while the precessing template bank cover the variables $M, q, s_1, \theta_1, s_{2z}, \iota$. For each template bank we report the range covered by each variable as well as the number of templates (size). The BBH variables not explicitly mentioned in the table are set to 0.

method, which builds a template bank through an iterative process, as discussed in Sec. 3.3. As the stochastic method is not suitable to cover high dimensional regions of the parameter space, such as those associated to precession, we generate the *precessing* template banks with our publicly available code `mbank` [410], introduced in Ch. 4. Our novel method relies on the random template placement to quickly cover the parameter space of interest, coupled with a normalizing flow model trained for an effective template sampling and ensures a very fast placement at the price of up to $\sim 50\%$ larger banks than those produced by the state-of-the-art stochastic approach.

As it is common, the non precessing template banks sample the variables

$$M, q, s_{1z}, s_{2z}$$

while both the precessing templates bank sample the variables

$$M, q, s_1, \theta_1, s_{2z}, \iota$$

where $M = m_1 + m_2$ denotes the total mass of the binary and the first spin is expressed in spherical coordinates:

$$s_x = s \sin \theta \quad (6.7)$$

$$s_y = 0 \quad (6.8)$$

$$s_z = s \cos \theta. \quad (6.9)$$

In our template bank, precession is only encoded into s_x ,² while all the other in-plane spin components s_{1y}, s_{2x}, s_{2y} are set to zero. Note that we also set the reference phase³ $\phi = 0$, despite the fact that the search statistic Eq. (6.2) does not explicitly maximizes over ϕ . The reasons for this choice were already discussed in Sec. 4.5.1.

To generate the precessing template banks, we use the frequency domain approximant `IMRPhenomXP` [47], while for the aligned-spin banks we use `IMRPhenomD` [157]. In all cases, we consider frequencies between $f_{\min} = 15$ Hz and $f_{\max} = 1024$ Hz and we use the Advanced LIGO O4 Design Power Spectral Density (with 190 Mpc range) [425].

To validate a template bank, we select a number of test signals (also called injections) and for each signal $s(\theta)$, characterized by θ , we compute the *fitting factor* FF Eq. (3.45), with the match \mathcal{M} computed with Eq. (4.1). The precessing injections are sampled from the relevant mass space using a probability distribution uniform in m_1 and q . The spins' magnitudes are sampled so that the fourth power of spin is

²Here s_x plays the role of the precession spin parameter χ_P .

³Note that the reference phase ϕ has not to be confused with the phase Φ_T of a trigger. While the first is a physical property of the system, the latter corresponds to the angle between the real and the imaginary part of the complex SNR timeseries $z(t)$ and it depends also on the location and orientation of the detector.

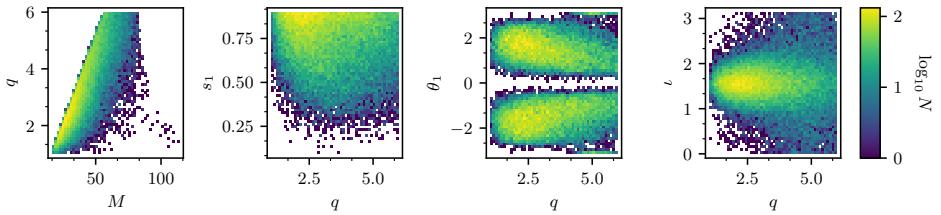


Figure 6.1: Binned templates of the “Low q ” precessing template bank. For each equal size bin, we color code the logarithmic number of templates so that the color is also a measure of the template density. In the different panels, we consider the variables $M, q, s_1, \theta_1, \iota$. The bank was first introduced in [1].

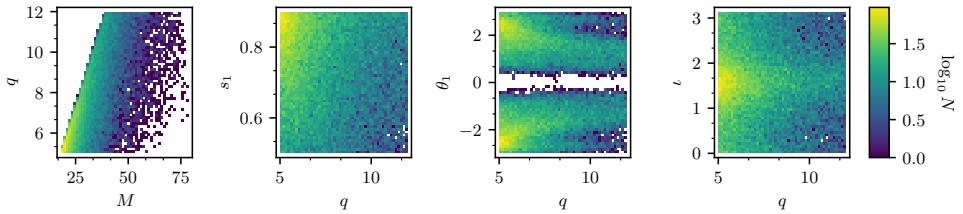


Figure 6.2: Binned templates of the “High q ” precessing template bank. For each equal size bin, we color code the logarithmic number of templates so that the color is also a measure of the template density. In the different panels, we consider the variables $M, q, s_1, \theta_1, \iota$.

uniformly distributed in $[0, 1]$. While this distribution of spins is clearly unphysical, it was chosen to make sure that a large fraction of the injected signals had a large precessing spin: in this way, approximately 70% of the injected signals have $\chi_P > 0.5$. The spin directions, the sky location and the binary orientation are all isotropically distributed over the solid angle. The reference phase is uniformly sampled in $[0, 2\pi]$. We generate a set of 21947 precessing injections, with a 30 seconds spacing. This will be used both for validating the banks and for the runs described in Sec. 6.4. To validate the aligned-spin banks, we use the *same* 21947 precessing injections as before, with the in-plane spin components set to zero (i.e. $s_{1x} = s_{1y} = s_{2x} = s_{2y} = 0$).

A template bank for low mass ratio systems The precessing “Low q ” template bank was introduced in [1, Sec. V]. It is a large bank covering systems with component masses $m_1, m_2 \in [8, 70]\text{M}_\odot$ with mass ratio limited into $q \in [1, 6]$. With a target minimal match of 0.97, the template bank has 1.6 million templates. We also generated an aligned-spin bank in the same mass region, consisting of 1.4×10^4 templates. A summary of the two “Low q ” banks is reported in Tab. 6.1.

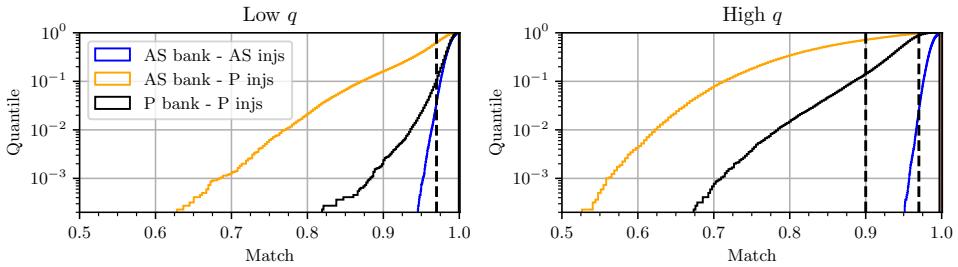


Figure 6.3: Fitting factor study for the four template banks considered in this work. For each given value of the match, we report the fraction of injections with fitting factor lower than that value. The two aligned-spin banks (AS) and the two precessing banks (P) are tested against the relevant injection set of fully precessing injections used for searching the data, with spins magnitudes reaching values up to 0.99. To evaluate their ability to cover the non-precessing space, the two AS template banks are also tested against a set of aligned-spin injections. The composition of the injection sets is described in the text. In the left panel, we report the results concerning the “Low q ” parameter space, while the right refers to the “High q ” parameter space. The vertical dashed lines mark the target minimal match of the template banks: 0.9 for the “High q ” precessing bank and 0.97 for the others.

In Fig. 6.1 we plot the template density of the precessing bank, while in Fig. 6.3 we present the results of an injection study. We note that both the precessing and the aligned-spin banks have an acceptable performance with only $\sim 10\%$ of the test signals having a match below the 0.97 threshold. The aligned-spin bank recovers $\sim 15\%$ of the precessing injections with a fitting factor below 0.9. Such a “small” fraction of signals below a given match threshold is in line with what is achieved by other template banks (see e.g. Fig. 4.3 and Fig. 4.9) and it suggests that our 0.97 aligned-spin bank approximates a 0.9 target minimal match bank for precessing signal. This remark will be needed to explain some of the search results.

As discussed in [1], the precessing bank provides a very poor coverage in the “low q , low M ”, characterized by $q \lesssim 1.2$ and $M \lesssim 20 M_{\odot}$. For this reason, we exclude such region from our searches (and also from our injection studies)⁴: systems in this region are expected to show little or no effect from precession, hence this region is of scarce interest for our purpose of quantifying the benefit of using precessing templates. Finally, we note in Fig. 6.1 that two small regions at $\theta = \pm\pi$ and $q \simeq 5$, corresponding to the aligned-spin limit, are unphysically populated by many templates. This is a feature of the metric approximation, probably due to the approximant in

⁴To address the issue of poor coverage, an effective strategy could involve using aligned-spin templates to cover the “low q , low M ” region.

use, and was also discussed in [1].

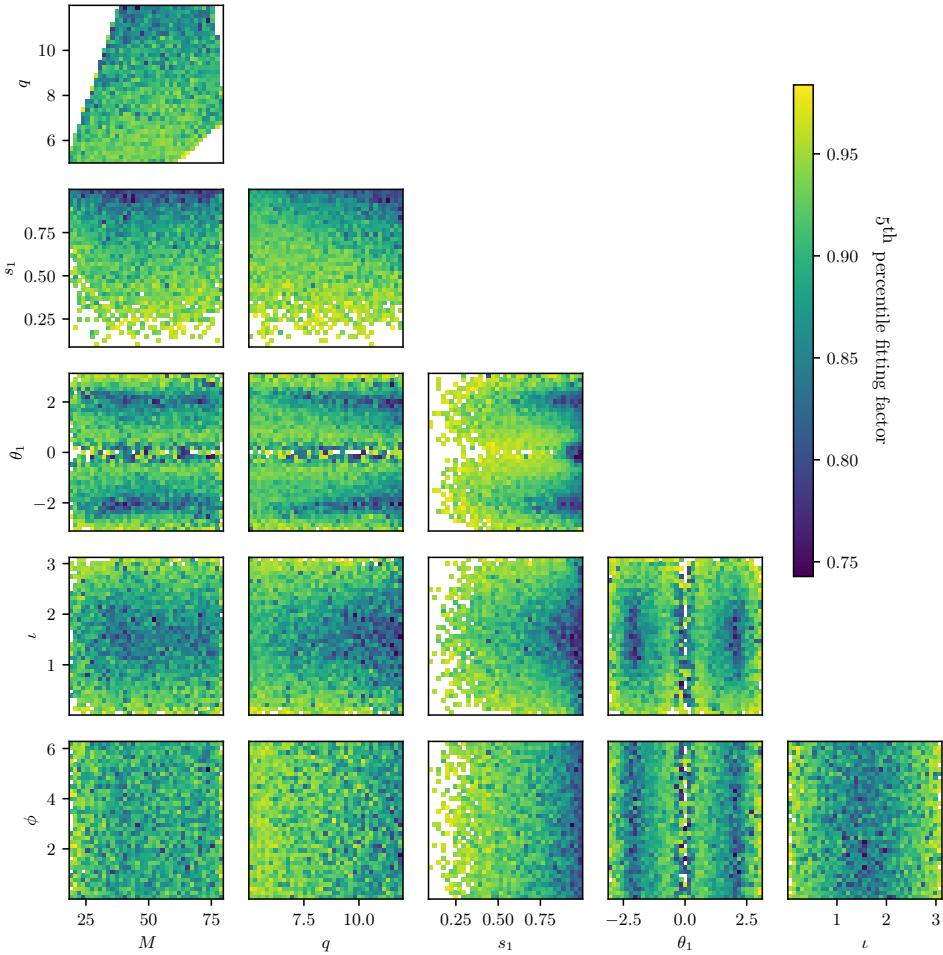


Figure 6.4: Fitting factor of the “High q ” precessing template bank, as a function of total mass M , mass ratio q , tilt θ_1 and magnitude s_1 of the first BH, inclination angle ι and reference phase ϕ . For each bin, we report the 5th percentile of the fitting factor distribution. Note that while the fitting factor study covers values of primary spin between 0 and 0.99, the template bank only gathers templates with $s_1 \in [0.5, 0.9]$.

A template bank for high mass ratio systems To target heavily asymmetric precessing systems, we construct the precessing “High q ” template bank. The bank targets systems with mass ratio $q \in [5, 12]$, with primary mass $m_1 \in [15, 70]M_\odot$ and secondary mass $m_2 \in [3, 10]M_\odot$. For the precessing bank, we limit the magnitude of the

primary spin s_1 to the range $[0.5, 0.9]$. As shown in the fitting factor study Fig. 6.3, our choice does not impact the bank injection recovery, in the low s_1 region excluded by the bank. As part of our development, we also generated a bank covering the entire spin range. However, this additional bank, consisting of 3×10^5 more templates, did not demonstrate improved coverage of the low s_1 region.

Given the vast size of the parameter space, employing a template bank with a target minimal match of 0.97 would entail dealing with an unwieldy number of templates, reaching as high as 30 million⁵. As modern pipelines are unable to handle banks of similar size, we have pragmatically set a target minimal match requirement of 0.9 for the precessing bank. With the bank now comprising 2.3 million templates, this adjustment helps manage the computational load while still ensuring a reasonable level of template coverage in the search for GW signals.

We also generate an aligned-spin bank, covering the same mass range: with a target minimal match of 0.97, the aligned-spin bank has 3.9×10^4 templates. It is noteworthy that the aligned-spin bank is two orders of magnitude smaller than the precessing bank. Consequently in this case, a precessing search is nearly two orders of magnitude more computationally demanding than its aligned-spin counterpart. Other details of the two “High q ” are reported in Tab. 6.1. Fig. 6.2 shows the template density of the precessing bank, where we observe that the bulk of the templates are placed in the high q , high s_1 region, particularly for systems with $\theta_1 \simeq \pm\pi/2$, as seen with an inclination $\iota \simeq \pi/2$. This distribution is consistent with the fact that signals in such regions exhibit the largest precession content. In Fig. 6.3, we report the results of an injection study for the two banks. Note that some of these injections lie outside the precessing template bank, for which $s_1 \in [0.5, 0.9]$. Nevertheless, our results show that both the precessing and “aligned-spin” bank are able to satisfy the relevant minimal match requirements of 0.9 and 0.97 respectively.

In Fig. 6.4, we study the dependence of the fitting factor of the precessing template bank on the several parameters characterizing a template. We note that the fitting factor is pretty stable across different regions of the parameter space, confirming the quality of our template bank. As the template bank only covers values of s_1 up to 0.9, we expect and observe a drop in accuracy for large primary spin values, which lie outside the template bank. The reduced accuracy at high spins seems to be limited only to edge-on systems (i.e. $\iota \simeq \pi/2$) with large values of χ_P (i.e. $\theta \simeq \pm\pi/2$), where precessing effects are the strongest. Finally, we see that the coverage is rather uniform in ϕ . This justifies our choice of not including the reference phase as part of the variables characterizing the precessing template banks.

⁵This number comes from an attempt we made to generate such bank with `mbank`. A similar number appears in [406], although referring to a template bank covering a slightly lower mass range.

For both the precessing and the aligned-spin bank, only $\sim 10\%$ of the injections are below the target minimal match threshold of 0.9 and 0.97 respectively. This is consistent with the results obtained in the “Low q ” case. With only $\sim 20\%$ of the precessing injections recovered with a match higher than 0.9, the aligned-spin bank struggles to cover the precessing parameter space satisfactorily. Note that in the “Low q ” case, the aligned-spin template bank was able to recover $\sim 85\%$ of the precessing injections with a match higher than 0.9.

6.3.1 SVD Compression

As opposed to the large majority of other pipelines, the GstLAL pipeline performs the matched filtering in time domain. Since computing the cross-correlation between two timeseries in time domain is notoriously more expensive than in frequency domain⁶, a Singular Value Decomposition (SVD) based template bank compression [77, 379, 444] scheme has been implemented to mitigate the cost of filtering.

As already mentioned in Sec. 3.4.1, templates in the banks are grouped together in bins of $\mathcal{O}(500)$ templates. Each template is then divided in different time slices, each with a different sample rate. For each time slice of each bin an SVD decomposition of the templates is performed. The data are then filtered only using a small subset of the SVD basis and the SNR timeseries for each template is reconstructed by “inverting” the SVD. Filtering the data using the SVD basis results in a reduction of one or two orders of magnitude of the number of filters. Clearly, the bins for the SVD must be carefully chosen: heuristically, each bin should gather very similar waveforms, so that only a small number of SVD basis are required to faithfully represent the templates.

The bins are constructed based on two quantities α^1, α^2 . As a first step, the template bank is split into N_{groups} bins based on the α^2 values. Second, groups of N_{tmp} templates are gathered together based on their α^1 value to form an SVD bin. Several choices have been implemented for the two sorting variables (α^1, α^2) . Common choices for α^1 are the template duration or the chirp mass \mathcal{M} :

$$\mathcal{M} = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}. \quad (6.10)$$

A typical choice for α^2 is the effective spin parameter $\chi_{\text{eff}} = \frac{m_1 s_{1z} + m_2 s_{2z}}{m_1 + m_2}$. As in the case of precessing templates, χ_{eff} does not capture the effect of precession, the effective precession spin parameter χ_P can be a good option for α^2 instead of χ_{eff} . In [380], the authors suggests to use the variables μ^1 and μ^2 , first introduced in [381], and

⁶Indeed, for a timeseries of D points, computing the correlation in time domain has a computational cost of $\mathcal{O}(D^2)$ while in frequency domain the same operation amounts to computing a Fast Fourier transform and has a cost of $\mathcal{O}(D \log D)$.

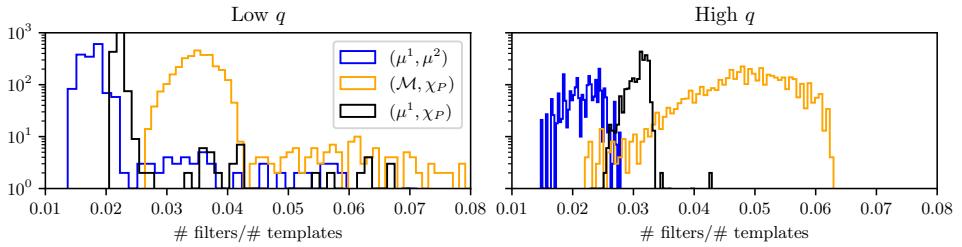


Figure 6.5: Compression achieved by the Singular Value Decomposition of the templates of the two precessing template banks considered in this work. For each bin, the compression is measured by the ratio between the number of filters and the number of templates. The number of filters is obtained by averaging, across different time slices, the number of SVD basis. The average is weighted by the sampling rate of each time slice. For different choices of the sorting variables (α^1, α^2) used to construct the bank split, we report the histogram of the compression achieved for different bins.

they achieve very high compression. The use of μ^1 and μ^2 is motivated by the analysis in [381], which shows that μ^1 and μ^2 are the two best constrained quantities by parameter estimation. Mathematically they are expressed as a linear combination of the first three coefficients Ψ^0, Ψ^2, Ψ^3 of the PN expansion for the frequency domain phase $\varphi^{\text{PN}}(f)$ of the waveform:

$$\mu^1 = 0.974\Psi^0 + 0.209\Psi^2 + 0.0840\Psi^3 \quad (6.11)$$

$$\mu^2 = -0.221\Psi^0 + 0.823\Psi^2 + 0.524\Psi^3 \quad (6.12)$$

The coefficients Ψ^0, Ψ^2, Ψ^3 depend on the two BH masses and on the z components of the two spins as well as on the minimum frequency considered for the template. The interested reader can find their expression in [380, Eqs. (9-11)].

In Fig. 6.5 we investigate how the compression provided by the SVD varies for the two template banks and for different choices of sorting quantities. The study aims at finding the best sorting quantities for maximal speed up. In our study, we set $N_{\text{tmp}} = 500$ and $N_{\text{groups}} = 20$.

Our results show that the sorting with variables μ^1 and μ^2 introduced in [380] provides the best speed up even in the precessing case. Regardless of the sorting variables considered, the “Low q ” precessing bank shows better compression than the “High q ” precessing bank. This is expected on the basis that the “High q ” bank gathers signal with higher precession content and more complex morphology. Therefore, a higher number of SVD basis is needed to faithfully reproduce the waveform.

A comparison with the results in [380] shows that, in the precessing case, the SVD compression is a factor of 2/3 worse than the aligned-spin case. As noted before, this

is expected due to the increased complexity of precessing waveforms. The higher number of SVD basis translates directly into a larger computational cost and it needs to be taken into account when allocating the computational resources for a precessing search.

6.4 Precessing Searches Results

Using the four template banks described in the previous Sec. 6.3, we run four searches on the publicly available LIGO and Virgo data [84] obtained through GWOSC [445], taken during the third observing run O3 between GPS times 1259423400 and 1260081799. As described above, we generate two sets of precessing injections, one for the “Low q ” banks and another for the “High q ” banks. This choice allows us to directly compare the results of a precessing search with those of an aligned-spin search.

When assessing the sensitivity of a precessing search compared to an aligned-spin search, two competing effects come into play. On one hand, the precessing search enhances the recovered SNR and lowers the ξ^2 , leading to improved sensitivity. On the other hand, the larger bank’s size results in an increased number of background triggers (i.e., false alarms), which can potentially downgrade the significance of any candidate event. Determining whether the improvement in recovered SNR outweighs the increased background triggers, and thus whether a precessing search indeed provides enhanced sensitivity will be the focal point of this section.

In what follows, we evaluate the improvement brought by a precessing search using three different figure of merits: recovered SNR, results from the ξ^2 test and, finally, increase in search sensitivity. To study the sensitivity we used the same injection set introduced for bank validation in Sec. 6.3, with injections sampled uniformly in m_1 and mass ratio inside the mass ranges covered by each template bank.

6.4.1 Recovered SNR

In Fig. 6.6, we present a histogram for each injection, illustrating the discrepancy ΔSNR between the SNR recovered by the precessing search and the SNR recovered by the aligned-spin search. The SNR difference is divided by the nominal SNR of the injection and we the data refers both to the “Low q ” and “High q ” cases. This analysis offers a valuable metric to evaluate the improvement brought by filtering the data with precessing templates compared to aligned-spin ones.

A certain degree of scattering of the discrepancies in recovered SNR (ΔSNR) is expected due to random noise fluctuations. This is especially true in the low SNR regime, where even a small SNR discrepancy can lead to a large fractional SNR variation. This remark is useful to explain the counterintuitive observation in Fig. 6.6 of a

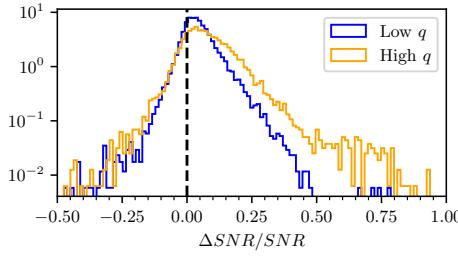


Figure 6.6: Histogram with the fractional SNR gain of a precessing search over its aligned-spin counterpart. For the two cases “Low q ” and “High q ”, we report the discrepancy $\Delta\text{SNR} = \text{SNR}_{\text{precessing}} - \text{SNR}_{\text{aligned-spin}}$ of the SNR recovered by the precessing and the aligned-spin, normalized by the SNR of each injection. The injections are described in Sec. 6.3 and the data refer to the LIGO Livingston detector.

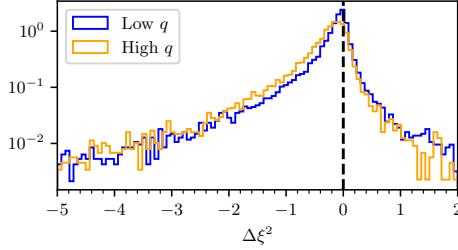


Figure 6.7: Histogram with the difference between the ξ^2 measured by a precessing search and its aligned-spin counterpart. For the two cases “Low q ” and “High q ”, we report the discrepancy $\Delta\xi^2 = \xi^2_{\text{precessing}} - \xi^2_{\text{aligned-spin}}$, evaluated on the set of precessing injections described in Sec. 6.3. Data refer to the LIGO Livingston detector.

tail of large negative SNR variations, for which the precessing triggers are recovered with lower SNR than their aligned-spin counterpart. Nevertheless, our analysis reveals a systematic trend: searches utilizing precessing templates consistently recover a larger fraction of the injected SNR compared to those using aligned-spin templates. This improvement can be substantial, with precessing searches recovering up to 75% more SNR than their aligned-spin counterparts. The systematic increase in recovered SNR fraction with precessing templates aligns with expectations, given the closer match of precessing templates to injected waveforms shown in Fig. 6.3. Whether the improved SNR recovery translates into a larger search sensitivity will be assessed in the next sections.

Finally, we observe that the improvement brought by the “High q ” precessing search is larger than in the “Low q ” case. This is consistent with our results in Fig. 6.3,

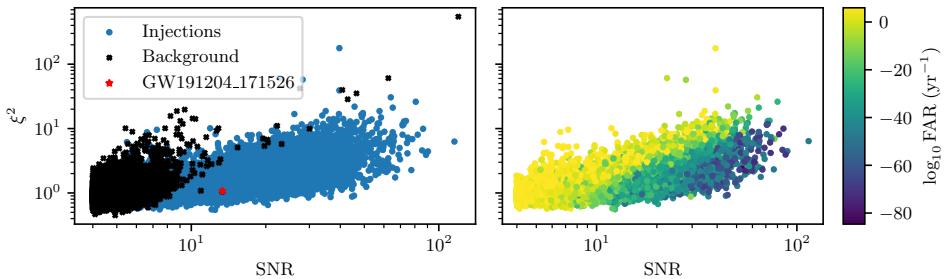


Figure 6.8: “SNR- ξ^2 ” plot for the “High q ” aligned-spin search. On the left panel, we report the SNR and ξ^2 of the background triggers, not associated to any injection, and the injection triggers, corresponding to an injection. On the right panel, we only report the injection triggers, colored by the logarithm of the FAR assigned by the pipeline. Data refer to the LIGO Livingston detector. The red star refers to the trigger produced by the GW event GW191204_171526 [18], detected by our search with $FAR < \frac{1}{1000} \text{ yr}^{-1}$.

which show that the aligned-spin “High q ” template bank exhibits poorer performance than the “Low q ” one, in recovering precessing injections. This is a direct consequence of the effects of precession being more prominent for asymmetric systems.

6.4.2 ξ^2 Test Results

We present in this section results from the ξ^2 consistency test. We will focus on the results of the “High q ” runs, comparing results from aligned-spin and precessing template banks. First of all, in Fig. 6.7 we report the difference between the ξ^2 values measured on the injected signals by the precessing and by the aligned-spin search. We note that in most cases, the ξ^2 measured by the precessing pipeline is lower than in the aligned-spin case. This is because the precessing templates are more similar to the injected signals, thus providing a more accurate prediction of the SNR timeseries and consequently yielding lower ξ^2 values. Moreover, this demonstrates the benefits of the updated ξ^2 expression in Eq. 6.5 over the standard expression Eq. 5.5.

In Fig. 6.8 and Fig. 6.9, we report the values of SNR and ξ^2 for the aligned-spin and precessing search respectively. The values refer to the LIGO-Livingston interferometer (L1). As clear from the left panel of Fig. 6.9, for the precessing run, the ξ^2 is able to separate real signals from the background for SNR higher than $\simeq 9$. The result is consistent with the standard behaviour on a set of aligned-spin injections [81].

In the Figs. 6.8-6.9 we see that both runs, aligned-spin and precessing, were able to detect the GW event GW191204_171526 [18]: the associated trigger is well separated

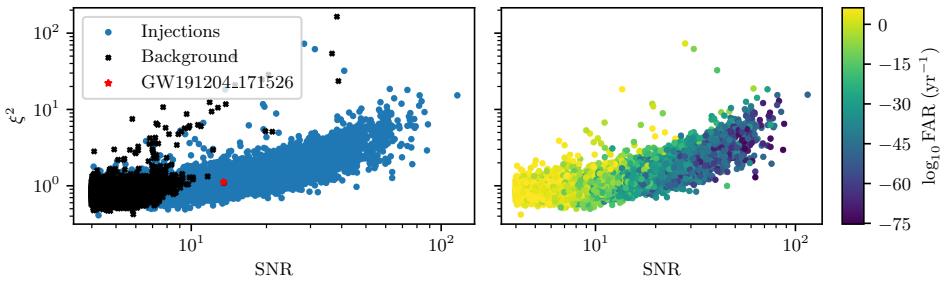


Figure 6.9: “SNR- ξ^2 ” plot the “High q ” precessing search. On the left panel, we report the SNR and ξ^2 of the background triggers, not associated to any injection, and the injection triggers, corresponding to an injection. On the right panel, we only report the injection triggers, colored by the logarithm of the FAR assigned by the pipeline. Data refer to the LIGO Livingston detector. The red star refers to the trigger produced by the GW event GW191204_171526 [18], detected by our search with $FAR < \frac{1}{1000} \text{ yr}^{-1}$.

from the background and hence was detected in both cases with a $FAR < \frac{1}{1000} \text{ yr}^{-1}$. In the figures, one can also see that the two searches record a handful of background triggers with high SNR and low ξ^2 , with the potential of contaminating the background of the searches. A close followup of these points revealed that such triggers were mostly caused by pairs of glitches happening very close in time to each other. This behaviour could possibly emulate the amplitude modulation typical of a precessing signal. It is worth noting that all of those potentially spurious candidates had a very high FAR, as a result of other terms of the ranking Eq. (3.54) downweighting their significance.

By looking at the right panel of Fig. 6.9, we note that the injection coincidences with $\text{SNR} \gtrsim 8$ are mostly ranked with a low FAR, meaning that they are mostly recovered. In the region characterized by $\text{SNR} \simeq 10$ and $\xi^2 \simeq 2$, the FAR slightly drops, suggesting that the ξ^2 test loses its effectiveness. This can be probably improved by using a denser template bank (at a much larger computational cost), by implementing the unapproximated ξ_{sym}^2 test or by improving the ranking statistic model.

In Fig. 6.8, it is evident that, in the case of the aligned-spin bank, the ξ^2 values for injection signals are considerably larger than in the precessing case. Clearly, an aligned-spin template is not able to model accurately the response of a heavily precessing signal, thus measuring a high ξ^2 value. Consequently, injection and background triggers are not well-separated in the $\text{SNR} - \xi^2$ plane and the ξ^2 has dramatically downgraded performance.

It is perhaps surprising that the reduced sensitivity of the aligned-spin search is not caused by a poor SNR recovery but rather by the failure of ξ^2 signal consis-

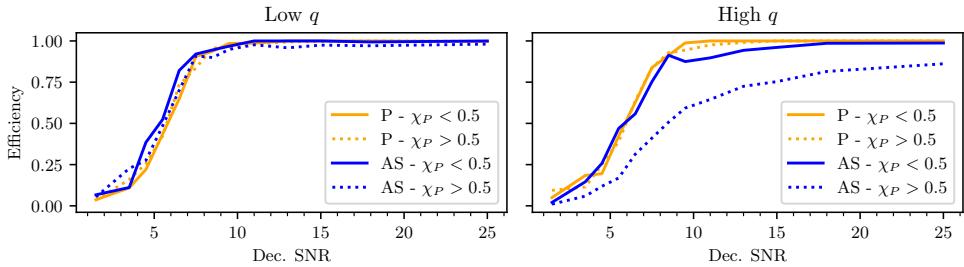


Figure 6.10: Efficiency of the BBH searches as a function of the loudest SNR measured by the three detectors, also called the “Decisive SNR”. The efficiency is defined as the fraction of signal detected at constant $FAR = 1/10$ years. In each panel, we show the efficiency pertaining the aligned-spin (AS) and the precessing (P) searches, reported separately for $\chi_P > 0.5$ and $\chi_P < 0.5$. The left panel refers to the “Low q ” search, while the right panel to the “High q ” search.

tency test. In other words, the sensitivity loss does not arise from using the “wrong” aligned-spin templates to filter precessing signals but rather from incorrectly labeling astrophysical triggers as non-Gaussian noise artifacts. This is manifest in the right panel of Fig. 6.8, where many injections recovered with a considerably large SNR present very large ξ^2 values, and are consequently assigned a high FAR. This observation aligns closely with previous literature on the topic, most notably [366, 382], while extending these well-established results to the precessing and/or HM case. Furthermore, it could significantly influence the design of future large-scale searches for precessing signals, as we will explore in Sec. 6.5.

An analysis of the “Low q ” runs (not reported here) shows similar conclusions, even though the downranking of injection triggers due to large ξ^2 has less impact on the search’s sensitivity, as will be discussed in the next section.

6.4.3 Sensitivity Improvement

We finally measure the improvement in the search sensitivity brought by the use of a precessing bank both in the “Low q ” and in the “High q ” case. As standard [82, 446], the sensitivity of a search is measured by spacetime volume (VT), corresponding to the volume reached by a search multiplied by the observing time T :

$$VT = T \int_0^\infty dV_c \epsilon(D_c, FAR) \quad (6.13)$$

where D_c is the comoving distance and $dV_c = dD_c 4\pi D_c^2$ is the infinitesimal comoving volume [447]. The quantity $\epsilon(D_c, FAR)$ is the efficiency of a search, defined as

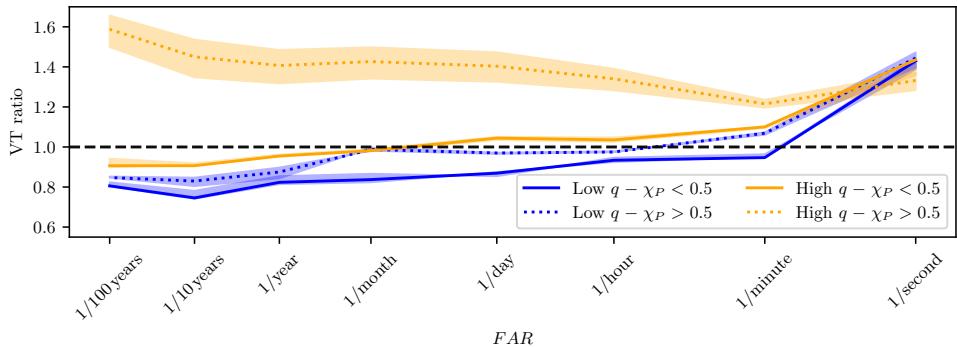


Figure 6.11: Sensitive volume increase as a function of the FAR detection threshold. The increase in sensitive volume is computed in terms of the VT ratio, measured between the VT of a precessing search and the VT of an aligned-spin searches. The two blue lines refer to the “Low q ” region, while the blue lines correspond to “High q ”. A dashed line refers to the VT computed only from injections with $\chi_P > 0.5$, while a solid line corresponds to $\chi_P < 0.5$.

the fraction of injections recovered in an infinitesimal shell between D_c and $D_c + dD_c$ with at least a given FAR. Clearly, every estimation for VT heavily depends on the distribution of injection: this motivated the use of the same injection set for the pairs of runs. For our computations

In Fig. 6.10 we report the efficiency (evaluated at a $FAR = 1/10$ years) of each search as a function “Decisive SNR”, defined as the maximum between the three SNRs measured by each detector. In the “Low q ” search, we note that there is not a significant difference between the aligned-spin and the precessing case. Moreover, the efficiency does not depend on the value of χ_P . In this case, the improvement in the SNR recovery and in the ξ^2 values brought by a precessing search does not translate into an increased efficiency. In the “High q ” case, the precessing search has comparable performance with the “Low q ”. On the other hand, the aligned-spin search has poorer performance, especially in recovering signals with $\chi_P > 0.5$. Thus, in the region covered by the “High q ” search, the use of precessing templates brought a substantial improvement in the pipeline’s sensitivity.

In Fig. 6.11, we report the ratio $\frac{VT_p}{VT_{AS}}$ between the VT of a precessing search and the VT of the corresponding aligned-spin search. The ratio is plotted as a function of FAR and we report different measurement for the low and high χ_P regions. We note that for very high values of FAR, the precessing searches always perform better than the aligned-spin ones. This is a direct consequence of the improved SNR recovery offered by the precessing search (see also Sec. 6.4.1 and Fig. 6.6). The picture changes as soon as a lower FAR threshold is considered. Due to the increased background

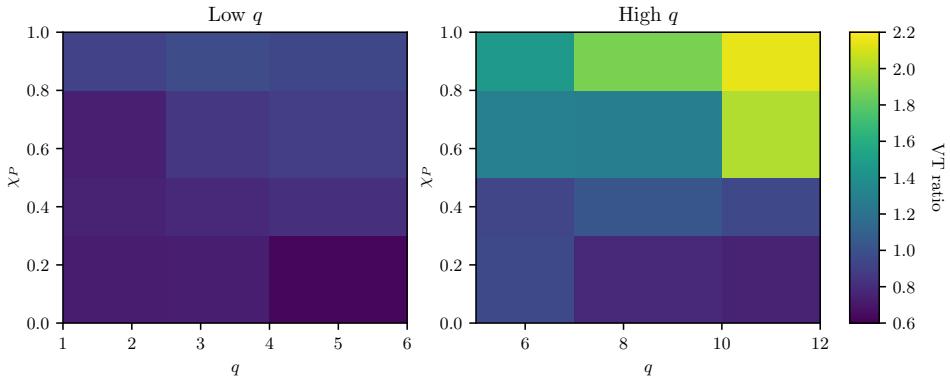


Figure 6.12: Sensitive volume increase for different bins in the $q - \chi_P$ space. The left panel refers to the “Low q ” region while the right panel to the “High q ” region. For each bin, we compute the VT ratio between a precessing spin search and the corresponding aligned-spin search. The VT ratio is encoded into the color scale and is computed at a $FAR = 1/10$ years.

caused by a large number of precessing templates, the “Low q ” precessing search suffers from a sensitivity loss of $\sim 20\%$ as compared to its aligned-spin counterpart. Indeed, even for highly precessing signals, with $\chi_P > 0.5$, the significance of many triggers is downranked by a large amount of false positives, due to a very large number of templates which can potentially lead to high SNR triggers. A more effective ranking statistic Eq. (3.54) can mitigate this issue by providing a more robust distinction between signal and noise. This approach was successfully applied in a search targeting precessing NSBH systems [406], where an updated ranking statistic was developed, improving upon its aligned-spin counterpart. Thanks to this enhancement, the authors of [406] achieved up to a $\sim 50\%$ sensitivity improvement compared to the aligned-spin ranking statistic. Notably, their search did not reveal a decrease in performance for low values of χ_P , unlike in our case. Future studies will be needed to understand the limitations of the current ranking statistic and to develop more powerful alternatives that are better suited for detecting precessing signals.

In the “High q ” scenario, we still observe a decrease in sensitivity for mildly precessing signals $\chi_P < 0.5$, even at low FAR. The observation above is still valid: the increased background downranks many potential candidates found by the precessing search. On the other hand, for $\chi_P > 0.5$ we observe a $\sim 50\%$ increase of the volume reached by the precessing search. The increase remains stable with the FAR. These considerations are consistent with our observations regarding the measured search efficiency Fig. 6.10.

It is remarkable, if not surprising, that in the “High q ” case an improvement in

sensitivity was obtained *despite* the fact that the precessing bank was generated with 0.9 target minimal match. Moreover, the “Low q ” aligned-spin bank, provides an excellent recovery of precessing injections, despite the fact that more than 10% of the injections considered have a match of 0.9 and lower. The two coupled observations seems to indicate that also future searches for precessing signals can employ 0.9 target minimal match template banks. This has the potential of reducing by orders of magnitude the computational cost as well as the background of the search.

In Fig. 6.12, the VT ratio is computed for different bins in the $q - \chi_P$ plane for a fixed FAR of 1/10 years. Results in in the left and right panels refer to the “Low q ” and “High q ” respectively and they corroborate what observed above. In the “Low q ” case, we see little variation of the sensitive volume across the space, and for low values of q and χ_P , we observe up to a 40% decrease in sensitivity due to the largely increased background. On the other hand, in the “High q ” scenario we note that VT ratio can be as high as 100%, depending on the values of q and χ_P . As already noted before, for low χ_P the precessing search does not bring any improvement. On the other, for $\chi_P > 0.5$, the sensitivity rapidly grows as χ_P and q grow, reaching a 120% gain in the most extreme case.

Our results suggests that targeting more asymmetric systems, such as a NSBH system, could potentially lead to an even larger sensitivity improvement. However as noted in [1], a search in such region struggles with an increased size of the parameter space, due to the increased length of the templates required. Indeed, a template bank in the NSBH region has hundreds of millions of templates and probably a different approach to the search might be needed. In this respect, the work presented in [406] is a promising avenue to target the precessing NSBH parameter space, due its lower computational cost.

6.5 Outline of Future Improvements

Up to this point, we have adapted the GstLAL pipeline to handle precessing templates by modifying the SNR computation routine and making minimal edits to the ξ^2 test, as outlined in Sec. 6.2. However, other components of the pipeline have remained unchanged. Although our modifications have successfully enhanced the pipeline sensitivity for precessing signals, there might still be room for improvement in the pipeline’s overall performance. Below, we summarize several potential directions for future developments.

Firstly, it’s important to acknowledge that the ξ^2 test implemented in this work represents an approximate version of the test proposed in [2] and discussed in Ch. 5. Implementing the full ξ^2 test would necessitate extensive changes to the pipeline,

which are beyond the scope of this work. In Ch. 5, an injection campaign in both Gaussian and real noise was conducted, suggesting that for signals with $\text{SNR} \lesssim 20$, the approximation employed in our implementation of ξ^2 does not introduce significant bias. However, for high SNR signals, the bias introduced by our approximation becomes more significant. We note that the study in Ch. 5 did not account for the discreteness of the template bank, which can potentially degrade the performance of the currently implemented ξ^2 test Eq. (6.5) in a real search scenario. Investigating this aspect is deferred to future work and is essential for determining whether modifying the ξ^2 test can lead to a substantial increase in the performance of a precessing search.

In our searches, we have not employed the “bank ξ^2 ”, introduced in [82]. The “bank ξ^2 ” is an alternative signal consistency test, which checks for a consistent response among several templates within the same SVD bin. It is used as an additional test to the “original” ξ^2 and it has been shown to provide an improvement in the high mass region. Future work should assess the performance of the “bank ξ^2 ” applied to precessing signals and possibly deploy it in a real search.

As previously discussed, the ranking statistic given by Eq. (3.54) may also need adaptation for precessing templates. The key LR term requiring modification is the “coherence” term, as highlighted in Sec. 6.2. In the precessing case, the measured phase of a trigger Φ_T loses the straightforward interpretation it has for an aligned-spin signal. Since this information has been crucial in producing the “coherence” term of the ranking statistic, it might not be directly suitable for a precessing template. Deriving an analytical expression that generalizes the “coherence” term to precession is likely to be challenging. In this context, utilizing an unsupervised Machine Learning model could prove to be an effective solution, allowing for a more flexible and data-driven adaptation of the ranking statistic to precessing templates.

Besides the “coherence” term, other factors of the ranking statistic retain the same physical meaning they have for an aligned-spin search. For this reason, in principle, they do not require modification. However, their predictive power depend on some hyperparameters, which were tuned by only considering aligned-spin searches. Tuning such hyperparameters on a set of precessing injections can also offer an increased performance for the search.

We observed in Sec. 6.4.2 that the primary loss in sensitivity by the aligned-spin “High q ” precessing search arises from the degraded performance of the ξ^2 test. This fact can have a profound impact on the search for precessing signals. Indeed, as the SNR recovery is not crucial for performance, it should be possible to deploy a template bank with a very low target minimal match (0.8, or even smaller). The reduced size of the resulting bank will have two favorable implications: (i) it reduces the com-

putational cost and (ii) it reduces the amount of false positives, possibly improving the sensitivity. Of course, a more sparse template bank will also downgrade the performance of the ξ^2 test and future work should explore the trade off between the benefits of a smaller template bank and the downsides of a less effective signal consistency test. With this strategy, it should become feasible to carry on a precessing search in the NSBH region of the parameter space, where the precession content of a signal is expected to be very strong.

Finally, it is important to stress that the changes we made to the pipeline are fully applicable to systems with HM content. Similar to the approach taken for precession in this work, future efforts will evaluate the pipeline’s performance in the presence of HM signals and assess potential sensitivity improvements. Once results from an HM search become available, the considerations outlined above regarding possible further pipeline upgrades remain relevant.

6.6 Final Remarks

In this chapter we describe the modifications made to the GstLAL pipeline to enable a systematic search for GW signals originating from precessing BBH systems. To filter the data, we incorporated the search statistic outlined in [183, 328, 427] by employing the “orthogonalized” template Eq. (6.3). To account for the effects of this modified template response, we made updates to the ξ^2 consistency test. The test now implemented is an approximation to the one proposed in [2] and discussed in Ch. 5.

To assess the effectiveness of our enhancements, we generated two large precessing template banks to search for precessing signals in two distinct regions of the BBH parameter space. These template banks were employed to search one week of publicly available data from the third observing run, spanning between GPS times 1259423400 s and 1260081799 s. The results obtained from these searches were then compared with two identical searches conducted using an aligned-spin template bank, which covered the same mass range.

As discussed in Sec. 6.4, searching the data with precessing templates results in a slight increase of a signal recovered SNR and to a substantial improvement in the ξ^2 test performance. However, this is achieved at the price of a large increase in the background triggers, due to a template bank being up to two orders of magnitude larger than in the corresponding aligned-spin case. We obtain a substantial improvement of the sensitivity of the pipeline only when searching strongly precessing systems, characterized by a high mass ratio $q \gtrsim 5$ and highly precessing spin $\chi_P \geq 0.5$. For less strongly precessing systems, we observe a slight decrease of the pipeline sensitivity, as the increased background removes any positive effect of the improved SNR.

recovery and of the more accurate signal consistency test.

In Sec. 6.5, we outlined a few possible future upgrades of the pipeline, including an updated ranking statistic Eq. (3.54) and the implementation of the “exact” signal consistency test introduced in Ch. 5. Moreover, the pipeline developed for precessing signals can be straightforwardly applied to the search of signals with HM content. Future work should aim at assessing the sensitivity brought by a similar search.

Our results lead to two significant conclusions that can inform the planning of future precessing searches. The first observation is technical: we find that employing a template bank with a target minimal match of 0.9 is sufficient to achieve a substantial sensitivity improvement for highly precessing signals. As discussed in Sec. 6.4.2, this is due to the fact that the primary sensitivity loss of an aligned-spin search targeting the precessing region is not due to a lack of recovered SNR but rather by a worsening of the ξ^2 value. The observation is crucial for reducing the size of the template bank, consequently lowering computational costs and the associated background in a search. The second observation pertains to the targeted parameter space. Our results suggest that, to attain any sensitivity improvement, a precessing search should focus solely on asymmetric and heavily precessing systems, characterized by $q \gtrsim 5$ and $\chi_P \gtrsim 0.5$. Current aligned-spin searches have limited sensitivity in such an extreme region, and indeed, signals in this range have not been detected by past aligned-spin searches.

Our work provides the community with the tools to search for precessing signals. Moreover, our investigations enable the community to select an appropriate parameter space to target with a precessing search, where a large sensitivity improvement, and possibly new detections, are within reach. In the next and final chapter, we leverage the capabilities of the precessing pipeline developed here to conduct a large scale search for precessing signals.

CHAPTER 7

Searching for Precessing Binaries in LIGO Data from the Third Observing Run

7.1 Introduction

As emphasized throughout this work, precession in BBH is of primary scientific interests with a profound impact on GW astronomy, ranging from BBH population studies [58–60, 62] to cosmology [27, 57]. As previously mentioned, precession has been observed as a statistical property of the BBH population [20, 21]. However, very few signals show decisive evidence for mis-aligned spins [53], as measured by χ_P . Interesting exceptions include GW190521 [54], GW191109 [55] and GW200129 [56]. It is worth noting however that analyses on GW200129 are made difficult by data quality issues [448, 449], which could possibly undermine any statement about the precessing nature of the system.

While the lack of strongly precessing signals might be due to their rarity, it might also be caused by the limited sensitivity of current searches to such extreme signals. If this is the case, it is important to explore thoroughly the BBH parameter space, using the most sensitive possible matched-filtering search pipeline. In view of this goal, in Ch. 6 we developed a precessing version of the GstLAL pipeline, we validated the results and we identified a region of the BBH parameter space where a sensitivity improvement of up to 120% is within reach.

In this chapter, we leverage the capabilities of the newly developed GstLAL precessing pipeline to search the LIGO-Virgo publicly available data from the third ob-

serving run O3 [84] for precessing BBH signals with component masses $m_1 \in [15, 70] \text{ M}_\odot$ and $m_2 \in [3, 10] \text{ M}_\odot$, with a mass ratio restricted to $q = m_1/m_2 \in [5, 12]$, using the “High q ” template bank introduced earlier. After briefly reviewing our search method in Sec. 7.2, we report the results of our search in Sec. 7.3. While we do not detect any novel signal, our results [293] allow us to place an upper limit on the astrophysical merger rate of a hypothetical subpopulation of asymmetric, heavily precessing signals, not detected by past searches. Results from this chapter are also presented in [4].

7.2 Set-up of the Search

To search for precessing signals, we employ the search technique introduced and implemented in [3], also described in Ch. 6, which offers up to 120% sensitive volume increase for very asymmetric highly precessing BBH systems, as compared to its aligned-spin counterpart. The sensitivity improvement is made possible thanks to three main features:

- *New metric template placement algorithm* [1]: by reducing the bank generation time by orders of magnitude, it allows for a fast identification of a region of the parameter space where a sensitivity improvement can be achieved using feasible number of templates.
- *Updated signal consistency test* [2]: it improves the robustness of the signal identification stage.
- *Implementation of a suitable search statistic* [328]: it allows for optimal SNR recovery of the candidates.

Our search employs the “High q ” 0.9 match¹ precessing template bank introduced in Ch. 6 with 2.3 millions of templates, targeting systems with component masses $m_1 \in [15, 70] \text{ M}_\odot$ and $m_2 \in [3, 10] \text{ M}_\odot$ and restricted to a mass ratio $q = m_1/m_2 \in [5, 12]$. Moreover, we only consider templates with $s_{1y} = s_{2x} = s_{2y} = 0$ and zero reference phase ϕ . The other spin components are set such that $s_1 \in [0.5, 0.9]$, $s_{2z} \in [-0.99, 0.99]$ and $\theta_1 \in [-\pi, \pi]$, with $\theta_1 = \arctan \frac{s_{1x}}{s_{1z}}$. As discussed in Ch. 4 and Ch. 6, such limitations of the spin degrees of freedom leverages the degeneracies in the spin parameter space to provide good injection recovery, while limiting the number of templates required to cover the space [1, 3]. We filter the data in a frequency range $f \in [15, 1024] \text{ Hz}$ with the waveform approximant IMRPhenomXP [47]. Note that for

¹The choice of using a 0.9 minimal match, instead of the standard 0.97, allows us to obtain a bank with a manageable size. As discussed in Sec. 6.4.1, it does not seem to have a large negative impact on the search.

our analysis we only consider *simple* precession, thus neglecting the effect of *transitional* precession². To compare the performance of our precessing search we use an aligned-spin template bank, also introduced in [3], covering the same mass range and with $s_{1z}, s_{2z} \in [-0.99, 0.99]$, employing 2.7×10^4 templates.

To estimate the search sensitivity and to place a limit on the merger rate of precessing signals (see Sec. 7.3.1), we sample a number of injections using the BBH population results inferred by the LIGO-Virgo-KAGRA (LVK) collaboration using GWTC-3 [21]. Our choice of relying on the population results from GWTC-3 poses a strong prior on the structure of the hypothetical subpopulation of asymmetric precessing system that we are searching for. To avoid this, we could have used a more “agnostic” population model, in order to avoid any preference for a population model developed only with aligned-spin searches. On the other hand, our choice makes sure that the rate upper limits from our search are comparable with those produced by the LVK. For this purpose, we use the best fit of the Binned Gaussian Process (BGP) population model [302], which allows us to easily sample BBH signals in the mass range of consideration. We sample the two spins isotropically, where the magnitude of each spin is constrained between 0.5 and 0.9. The sources are uniformly distributed in the volume between [30, 300] Mpc with uniformly sampled orientation, and generated with the IMRPhenomXP approximant.

For our search we use publicly available GW data [84] collected by the two LIGO observatories [11, 214, 450] during the third observing run (O3). Data collection periods were divided in two halves: the first half of O3 (O3a) spanned between 1 April 2019, 1500 UTC and 1 October 2019, 1500 UTC, while the second half (O3b) happened between 1 November 2019, 15:00 UTC and 27 March 2020, 17:00 UTC. Forced by limited computational resources, we did not use Virgo data taken in the same period, due to the smaller contribution to the sensitive volume given by Virgo.

7.3 Search Results

The search detected 30 events with $p_{\text{astro}} > 0.5$: they were all previously reported in the GW transient catalogs GWTC-2 [16, 17] and GWTC-3 [18] released by the LIGO-Virgo-KAGRA collaboration. The confident detections are reported in Tab. 7.1. In Fig. 7.1, we summarize the information about our detected GWTC signals, where we report all the detected events in GWTC-2 and GWTC-3. Each event not discovered

²As discussed in Sec. 1.3.3, transitional precession happens when the total spin vector \mathbf{S} is approximately opposite to the orbital angular momentum \mathbf{L} . In this case, the system does not have a rotation axis for the precession motion, which is usually given by the total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$. As a consequence, the orbital orientation will rapidly changes until a new configuration (where \mathbf{J} is non-zero) is reached.

	SNR	SNR _{GWTC}	FAR(1/yr)
GW190408_181802	14.29	14.6	$< 10^{-5}$
GW190412_053044	17.82	19.8	$< 10^{-5}$
GW190512_180714	11.88	12.7	$< 10^{-5}$
GW190513_205428	11.8	12.5	0.0008
GW190519_153544	12.69	15.9	0.0001
GW190521_074359	22.74	25.9	$< 10^{-5}$
GW190527_092055	8.79	8.0	0.1547
GW190706_222641	12.54	13.4	0.0221
GW190707_093326	12.68	13.1	$< 10^{-5}$
GW190708_232457	12.13	13.4	$< 10^{-5}$
GW190720_000836	10.55	10.9	0.0012
GW190727_060333	11.3	11.7	$< 10^{-5}$
GW190728_064510	11.92	13.1	$< 10^{-5}$
GW190814_211039	20.52	25.3	$< 10^{-5}$
GW190828_063405	15.93	16.5	$< 10^{-5}$
GW190828_065509	10.56	10.2	0.001
GW190915_235702	12.09	13.1	0.0015
GW190924_021846	12.87	12.0	$< 10^{-5}$
GW191109_010717	14.19	17.3	$< 10^{-5}$
GW191129_134029	13.14	13.1	$< 10^{-5}$
GW191204_171526	16.78	17.4	$< 10^{-5}$
GW191215_223052	10.36	11.2	0.0008
GW191216_213338	17.52	18.6	$< 10^{-5}$
GW191222_033537	10.91	12.5	0.0009
GW200128_022011	9.0	10.6	0.0059
GW200129_065458	25.8	26.8	$< 10^{-5}$
GW200224_222234	17.85	20.0	$< 10^{-5}$
GW200225_060421	12.1	12.5	0.0664
GW200311_115853	15.69	17.8	$< 10^{-5}$
GW200316_215756	9.96	10.3	0.0003

Table 7.1: Confident events detected by our search with $p_{\text{astro}} > 0.5$. For each event, we report the event name, together with the SNR and FAR. For reference, we also report the SNR_{GWTC} reported in the transient catalog.

by our search is represented by a cross, while a large dot, colored by the detection FAR, represents a detection. Note that we report events detected by our search up to a FAR = 1/day, hence some of them do not satisfy our detection threshold of $p_{\text{astro}} > 0.5$.

GPS time	SNR	FAR(1/yr)	p_{astro}
1252415231	8.84	2.367	0.06
1252465013	8.35	3.2005	0.05
1253504581	9.02	4.5486	0.03
1267433277	7.75	5.4287	0.04

Table 7.2: Candidates not reported in GWTC, detected by our search with $\text{FAR} < 6/\text{yr} = 1/2 \text{ months}$. For each candidate, we report the SNR, FAR and p_{astro} .

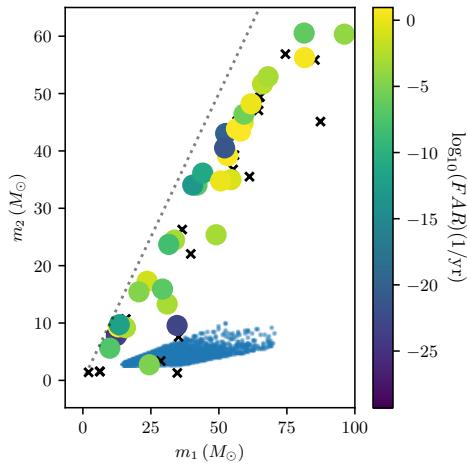


Figure 7.1: Median masses of the GWTC candidates as computed with parameter estimation. The GWTC candidates detected by our search are denoted by large circles, colored according to their FAR. A black cross marks the GWTC events that our search missed. The blue dots represent the templates used in our search bank.

As shown in Fig. 7.1 and as further confirmed by parameter estimation studies [16, 18], the GWTC events detected by our search lie mostly *outside* the region covered by our template bank. We observe that these events typically have a higher total mass than our templates, meaning they are shorter in duration. Consequently, we conclude that the sensitivity of our search extends to higher masses beyond our intended target. This is likely due to the large number of relatively long-duration templates used, which can easily match short-duration events, albeit with sub-optimal SNR and ξ^2 values. The fact that the detected events are outside the target region also explains the observation in Tab. 7.1 that our search measures a systematically lower SNR than the GWTC searches. Finally, we highlight that our search was able to recover with high confidence the three individual events previously mentioned

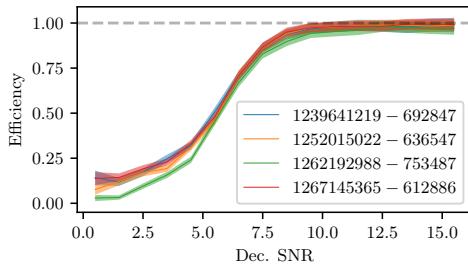


Figure 7.2: Efficiency of the search as a function of the decisive SNR, i.e. the maximum injected SNR between the two detectors. The efficiency is measured on four different chunks with GPS start time and length reported on the caption separated by a dash.

(GW190521, GW191109, and GW200129), for which spin misalignment measurements had been claimed by other works [54–56].

In Fig. 7.1 we note that two events, GW191113_071753 and GW200210_092254, were missed by our search, even if they lie close to our target region. The first one GW191113_071753 was not recorded by the GstLAL pipeline even in the offline broadband search³ and for this reason it is not surprising that also our precessing search failed to detect it. On the other hand, the issue of why GW200210_092254 was missed requires more investigation. For the moment we can postulate that this is due to a combination of the signal being at the very edge of our template bank and of the low mass of the lighter object, resulting in a signal duration longer than most of the templates in the bank. Unlike the other short events detected outside the bank, the longer duration of this event makes it more challenging for a template to recover the entire SNR of the system.

Besides the known GW events, our search did not detect any novel GW candidate. In Tab. 7.2, we report a list of the candidates observed with $\text{FAR} < 2/\text{year}$, not included in a previous transient catalog.

In Fig. 7.2, we report the efficiency $\epsilon(\text{SNR}, \text{FAR})$ of our search for different bins of *decisive* SNR, defined as the maximum injected SNR between the two LIGO detectors. The efficiency amounts to the fraction of signals detected by the search and it is computed using the BGP injection set defined above. We perform injections only for two chunks of data in O3a and O3b respectively, and we report the results separately for each chunk. The efficiencies obtained are similar to those reported in [3], however for the chunk starting at GPS time 1262192988 s we observe a drop in efficiency of

³The event was found instead by the PyCBC [355–357] and Multi-Band Template Analysis (MBTA) [360, 361] pipelines, although with a low significance. Similar situations, where only a few pipelines find a given event, are very common for marginally significant events.

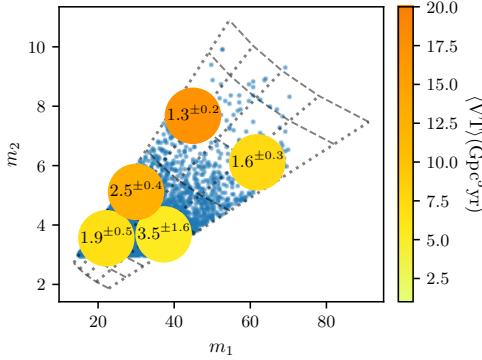


Figure 7.3: Search sensitivity computed over different regions of the parameter space. Each dot marks the center of a set of log-normally distributed injections and it is colored according to the sensitive spacetime volume measured using such injection set. The number superimposed denotes the ratio between the sensitive spacetime volume of our search and the one of an aligned-spins search targeting the same space.

unknown origin.

The efficiency is used to compute the averaged sensitive spacetime volume $\langle VT \rangle$ of the search, defined as:

$$\langle VT \rangle(\text{FAR}) = T \int_0^\infty dV_c \, \epsilon(\text{SNR}(D_c), \text{FAR}) \quad (7.1)$$

where dV_c is the comoving volume element at a given comoving distance D_c . In Fig. 7.3, we report the VT as a function of the parameter space of our search. The efficiency is averaged over the four chunks for which it was computed and it is assumed constant for all the observation time. We also report the $\langle VT \rangle$ improvement brought by our search as compared to an aligned-spin search. For each centroid, we employ $\sim 3 \times 10^4$ injections, log-normally distributed around the center with a variance of 0.1. All the sensitive volumes are estimated at a fiducial FAR = 1/yr.

7.3.1 Rate Upper Limits

The lack of newly detected signals allows to place an upper limit [293, 294, 392, 393] on the presence of a subpopulation of asymmetric, heavily precessing signals, not identified by other searches. In our analysis, we assume that such population is well described by the BGP injection set and, unlike standard works in BBH population analysis [20, 21], we do not attempt to constrain the distribution of BBH masses or spins.

After removing from consideration all the triggers associated to any GWTC event,

FAR = 1/yr
$\langle VT \rangle_{\text{O3a}} = 4.10_{-0.26}^{+0.25} \text{ Gpc}^3 \text{yr}$
$\langle VT \rangle_{\text{O3b}} = 3.2_{-1.5}^{+1.5} \text{ Gpc}^3 \text{yr}$
FAR = 2.36/yr
$\langle VT \rangle = 8.4_{-4.6}^{+1.5} \text{ Gpc}^3 \text{yr}$

Table 7.3: Sensitive spacetime volume estimated for the BGP population at a nominal FAR of 1/year and the FAR corresponding to the loudest event not reported in the GWTC. In the first case, we report the $\langle VT \rangle$ separately for O3a and O3b.

the loudest trigger has a $\text{FAR} = 2.367/\text{yr}$. We can then use our knowledge of the sensitive spacetime volume to place an upper limit $R_{90\%}$ (90% confidence interval) on the merger rate of the said population of heavily precessing signals. From this, we can compute the 90% confidence limit for the upper limit to the said population of heavily precessing signals,

$$R_{90\%} = \frac{2.3}{\langle VT \rangle} \quad (7.2)$$

where the sensitive spacetime volume is estimate using the FAR of the loudest event. Using the VT values for the BGP population reported in Tab. 7.3, we obtain an upper limit of

$$R_{90\%} = 0.28_{-0.04}^{+0.33} \text{ Gpc}^{-3} \text{yr}^{-1}.$$

The rate upper limit we obtain is consistent with the LVK rate estimates for our sub-population [21]: considering that our target population is only a small fraction of the total BBH population considered in the analysis, the LVK results imply a BBH merger rate of

$$R_{\text{LVK}} = 0.117 - 0.298 \text{ Gpc}^{-3} \text{yr}^{-1}$$

for our target region. Thus, our results do not provide any evidence to support the hypothesis of a population of asymmetric, heavily precessing signals that have not been detected by other searches. Searching the data taken during the first and second observing runs [83] will increase the surveyed sensitive spacetime volume, however it is unlikely that such increase will change the picture outlined above.

According to our results, the past analysis carried on by the LVK have not missed any heavily precessing asymmetric system, despite having a poor sensitivity towards those signals. Therefore, the population analysis based on the confirmed events of the GWTC [20, 21, 76] retain their validity in light of our findings, at least in the mass region target by our search. In particular, we are able to confirm that precession is a rare event in BBHs, detectable in $\sim 2\%$ of the cases with the current detector's sensitivity, as pointed out in [76].

7.4 Conclusion

In this chapter, we presented the results of a matched-filtering search on the interferometric data from the LIGO-Virgo-KAGRA third observing run. Our search targets asymmetric precessing BBHs, with a doubled sensitivity over such extreme systems: the search method is thoroughly discussed in [3]. We report the detection of 30 candidates, already announced in the LIGO-Virgo-KAGRA transient catalogs, but we do not find any new candidate. Our results allow us to place an upper limit on the merger rate of the binaries targeted by our search and, as our result is consistent with the rate estimates from the LVK, we can rule out the existence of a subpopulation of asymmetric precessing BBH not detected by past searches. Therefore, our results support the conclusions drawn from the existing catalogs: precessing systems are as “rare” as we thought they were before conducting the search.

Future work should focus on expanding the parameter space of the search, possibly targeting Neutron Star- Black Hole systems. This attempt will struggle with the huge size of the template bank required for this task, although see [406] which demonstrated that this is possible with only a 3x increase in template bank size.

CHAPTER 8

Closing Remarks

In this dissertation, we described the four years long journey to detect gravitational waves emitted by precessing binary black holes systems, also called precessing signals. Our work was motivated by the apparent rarity of detection of precessing signals [76], which raised the obvious question of why so few signals with clear imprint of precessions have been observed. Given the large scientific interest of precessing signals, with impacts on BBH formation mechanisms and cosmology, the question is very worthy of an answer. Therefore, to rule out the possibility that a few precessing signals were missed by the lack of sensitivity of current aligned-spin searches, we developed a technique to search for precessing signals, achieving up to 120% sensitivity increase over traditional searches. Thanks to our work, we were able to search for the first time for precessing signals in a small but interesting corner of the precessing parameter space and we were able to rule out the presence of precessing signals not detected by previous searches, thus confirming the results from the LIGO, Virgo and KAGRA collaboration.

The search results, which we presented in Ch. 7 and published in [4], were made possible thanks to three preparatory works: the first two addressed key challenges, while the third detailed and validated our end-to-end pipeline for precessing searches. The first of our work, presented in Ch. 4 and in [1], was dedicated to the daunting task of placing templates in a high dimensional space, such as the one in which precessing signals are embedded. This task was complicated by the strict requirements on the bank's effectualness as well as on its size. Building on solid theoretical research as well as on effective computational method, we were able to cast the prob-

lem of placing templates into the problem of sampling from a manifold. The sampling was elegantly performed with a normalizing flows, a machine learning model widely used to sample from complicated probability distributions. Thanks to our work, publicly available as a python package `mbank` [410], we are able to place *millions* of templates in a matter of hours: this is a significant improvement over previous methods, which required days of memory-intensive computations for generating banks. While primarily motivated by the search for precessing signals, the relevance of our work extends well beyond that. Indeed, it can be straightforwardly applied to generate template banks for eccentric signals or signals with imprints of HMs and it can find applications even outside the problem of template placing.

A second work, discussed in Ch. 5 and first presented in [2], was concerned to solve a technical, but important, issue for matched-filtering searches: what is the best signal consistency test to use for a precessing search? After the filtering stage, where the correlation between a template and the data is computed (SNR timeseries), the signal consistency test checks whether the measured SNR timeseries corresponds to the theoretical expectations. A robust signal consistency test is able to discern between astrophysical signals and short noise transient bursts happening in the detectors. While traditional tests were designed and tested only for aligned-spin systems, our study allowed us to introduce a novel generalized test ξ_{sym}^2 , which does not make any assumption about the nature of the signal. Moreover, we obtained a computationally convenient approximation ξ_{mix}^2 for the newly introduced test. Finally, we assessed the validity of the two newly introduced tests and we compared their performance with the traditional one. Our results suggested that for precessing signals it is crucial to introduce either the novel test or its approximation.

Our findings on the signal consistency test, together with our novel template placement method, allowed us to upgrade the GstLAL pipeline to effectively search for precessing signals. The new pipeline implements a suitable expression for the SNR timeseries computation, previously introduced in [328], and utilized ξ_{mix}^2 as a signal consistency test. In the work, presented in Ch. 6 and published in [3], we identified two interesting regions of the precessing parameter space, for both of which we performed a precessing search and compared it with an aligned-spin search over the same mass range. It is worth to stress that the identification of a region to search is already a computationally challenging task, as we have to balance between the desire of a search covering a wide region of the parameter space and the need of having a contained number of templates. We could have not achieved this delicate balance without the capabilities of our template placement code `mbank`. The two precessing searches have roughly 100x more templates than their aligned-spin counterparts. While this surely improves the SNR recovery and the effectiveness of the signal con-

sistency test, it increases the number of false alarms, potentially washing out any improvement. Our results indicate that the latter effect prevails in the case of system with a low mass ratio $q \lesssim 5$: in this scenario, our precessing search does not offer any improvement, over the traditional method. A different conclusion was drawn for the “High q ” case, covering systems with mass ratio $q \in [5, 12]$ and with primary mass $m_1 \in [15, 70]M_\odot$ and secondary mass $m_2 \in [3, 10]M_\odot$. In this region of the parameter space, the improvement in SNR and ξ^2 prevails and the precessing search offers up to 120% sensitivity increased for high values of mass ratios and effective precession spin parameter χ_P . Therefore we conclude that, in this “High q ” region, the sensitivity of traditional searches is very limited and a search with precessing templates is very much required. This paved the way for the first ever large scale search that we conducted on data from the LVK third observing run.

Too costly? The increase in sensitivity achieved by our precessing search comes at a heavy price. The two orders of magnitude increase in the size of the template bank translates almost directly into a two orders of magnitude increase of the computational cost of the search. Indeed, the cost of the search described in Ch. 7 can be quantified in half a million of CPU hours! While such cost is still manageable by modern computer clusters, we were only able to search a small region of the parameter space. A more extensive search for precessing signals, for instance covering also the interesting region of NSBH systems, can easily add a factor of ten on this cost, making similar searches unfeasible. Of course, this sets a large limitation to the breadth of precessing searches and it poses a threat to our scientific program of systematically searching for exotic signals. For this reason, a future generation of precessing searches must do better than the “brute force” approach we employed in this work. A very interesting example is presented in [406], where a proof-of-principle search for NSBH systems is developed with only a 3x increase in computational cost. The search, however, relies on an approximate expression for the precessing waveform and it is only valid for precessing signals composed by the dominant $\ell = |m| = 2$ multipole.

Based on our findings, we are able to suggest a different approach to speed up our precessing searches. This might lead to a promising search for systems which includes not only precession but also contributions from higher order multipoles, at a moderate computational cost. While this path has not been yet explored, we sketch below a few ideas that can help to focus the efforts of an interested researcher. A crucial takeaway from Ch. 6 is that the drop in performance of an aligned-spins search does not only arise from a low SNR recovery but also from a poor performance of the signal consistency test (see e.g. Fig. 6.8). In other words, even though an aligned-spin

template has a high correlation (SNR) with data containing a precessing astrophysical signal, the template response differs from the expected one, thus delivering a high ξ^2 value. This means that a search with aligned-spin templates will be likely launch a trigger in correspondence to an heavily precessing signal but this trigger will be assigned a low significance due to its high ξ^2 value and hence discarded. This suggests that a search targeting precessing signals can filter the data with aligned-spin templates and, in a second stage, recompute the SNR and ξ^2 values of the “promising” triggers, i.e. of triggers showing SNR values above a certain threshold. This second stage should aim at minimizing the ξ^2 values associated to each trigger with a set of precessing templates and advanced optimization techniques might be also used to efficiently find the template minimizing the ξ^2 .

This technique could save an enormous amount of resources, since the expensive filtering with a large number of precessing templates would only involve a small chunk of data. The advantages are manifest: with only a minimal increase in filtering cost, we can achieve a precessing search with largely enhanced sensitivity. Potentially, this approach could be easily extended to precessing systems with imprints of HMs and (why not?) imprints from an eccentric orbit. While this research project is yet to be fully explored, we believe that it has great potential and as such it constitutes one of the greatest “lessons learnt” from this dissertation.

Interesting times ahead! Regardless of the search techniques employed, the importance of a search that overcomes the limitations of the traditional ones is undeniable. The scientific impact of the detection of many precessing systems and/or systems with imprints from HMs is too wide to be neglected. This is also witnessed by the increasing attention given in the literature to the topic, with multiple groups competing to develop various state-of-the-art pipelines. Therefore, it is easy to predict that the next generation of searches for the coalescence of compact objects will incorporate an ever increasing number of different physical effects, hopefully leading to exciting new detections. This in turn can improve our understanding of many open problems in gravitational-wave astronomy. The present dissertation is a modest contribution to the wealth of different ideas and search techniques currently being developed by multiple groups worldwide. Nevertheless, we hope that the results from our search and the lessons learned during our journey will assist the community in generating new ideas or refining search methods, and in some small way, facilitate the new discoveries that lie ahead. We could not wish for a better outcome for the four years of effort we’ve dedicated to this work.

Public Summary

General relativity, black holes and gravitational waves Challenging our conventional understanding of concepts like time, energy, and space, the theory of general relativity is considered an accurate description of physical phenomena involving massive objects in the universe, such as planets or stars. Behind the intricacies of the mathematics, the theory predicts the existence of a class of a very dense star-like objects, called black holes (BH), not predicted by the classical Newtonian theory of gravity. A black hole is a region of the space where the gravity is so strong that nothing can possess enough speed to escape from its gravitational attraction. Therefore, anything entering a black hole, including light, will remain inside a black hole forever¹.

If black holes are the most famous prediction of the theory, they are not the only one. Indeed, the theory predicts also the existence of gravitational waves, arising as natural consequence of the finite speed of propagation of the gravitational interaction. A simple thought experiment can help to clarify this. Imagine to put a large mass in an empty space (in vacuum) and to probe with a small test mass its gravitational effect, which according to high school textbooks manifests itself as an attraction towards the large mass. Now imagine that (for whatever reason) at time t_0 the mass starts to move: what happens to the test mass? Or better, when does it feel the change in the attraction due to the fact that the mass have changed its position? According to the Newtonian theory of gravity, the interaction is instantaneous: as soon as the large mass moves, the small mass will feel the effect. While this feature of the theory was already looked with suspicion by Newton and his contemporaries, it took more than two hundreds years to establish that nothing can travel at “infinite” speed. For this reason, general relativity must predict that the test mass perceives a variation in gravity only at some delayed (retarded) time t_R . What happens between t_0 and t_R ? The gravitational interaction is propagating through the form of a gravitational wave! Of

¹This motivates their name: since no light is emitted or reflected, they appear as a purely black sphere to any observer “looking” at them.

course, this picture picture is heavily simplified and there are a lot of caveats on what you may or may not define as a gravitational wave. Nevertheless, the main physical idea still stands: a gravitational wave is a propagating gravitational interaction, pretty much like light is a propagating electromagnetic interaction.

Black holes binaries Once we establish that gravitational waves exist (note that this step is not trivial, as it took almost fifty years of scientific debate to establish their existence) we can ask ourselves: how are gravitational waves produced? Diving into all the mathematical details of the theory, provides a rather simple answer: gravitational waves are produced by accelerating objects. Of course, there are many accelerating objects in the universe, including a train leaving a station, but a particularly intriguing class of sources of gravitational waves comes from binaries of compact objects. Some definitions are in order: a compact object is a very dense stellar-like object, such as a neutron star or a black hole, while a compact binary is a pair of compact objects, which orbit around each other under the influence of the mutual gravitational attraction. In simple words, a compact binary is the familiar system of Earth and Sun, where both the Earth and Sun are replaced by heavier and denser objects. We call binary black hole a system formed by two black holes. Both the objects are in (quasi) uniform circular motion and, more importantly, they experience a centripetal acceleration: this means that they emit gravitational waves!

Due to their importance as GW emitters, the physics of binary black holes is widely studied in most of its details and indeed, very accurate predictions for the GW signal emitted by the system are known. Despite the mathematical intricacies, a few simple physical ideas arise. First of all, it is well known that, as the gravitational radiation carries energy, the two objects compensate for the energy loss by orbiting closer and closer to each other, until they eventually merge into a single object. Moreover, the spins (i.e. the rotation around a given axis) of the two objects play a role in the dynamics of the system, slowing down or accelerating the GW emission. Finally, if the two spins are mis-aligned with each other the system undergoes precession. This means that the orbital plane itself wobbles and changes orientation as the two black holes spiral towards each other, introducing unique signatures in the gravitational-wave signals.

The observation of the GW emitted by the coalescence of compact objects allows scientists to answer a huge variety of important scientific questions. First of all, one may wonder how many binaries exist in the universe, what are their properties and how they formed—this goes into the field of the so called population studies. Another set of questions concerns the validity of the theory of general relativity to describe such complicated binary systems—this is the field of tests of general relativity.

Moreover, it is possible to constrain the properties of dense matter, by observing the merger of neutron stars and the complex phenomenology attached. Finally, as the sources are placed at very large distances, we can use them to study the structure of our universe, posing constraints on its expansion rate and on its history, contributing to the field of cosmology.

Detecting gravitational waves from black hole binaries The list above just gave a hint of the scientific potential of the observing compact binary systems and it is no surprise that billions of euros were allocated to build large instruments dedicated to the detection of gravitational waves, such as LIGO, Virgo or KAGRA. Such instruments are designed to measure a tiny variation of the distance between two test masses, which is the effect of a gravitational wave on the spacetime predicted by general relativity. Such a precise measurement can only be achieved with very large interferometers, by analyzing the interference pattern of two lasers shot in the two kilometers-long arms. The interferometer outputs a time dependent estimation of the differential arm length (i.e. the difference between the length of the two arms), which is closely related to any incoming gravitational wave.

With the theory of compact binaries at hand and with sophisticated instruments operating to measure the gravitational waves, it would seem that the hardest part of the work is done. Unfortunately, most of the detector's output is noise. Therefore, without proper data analysis techniques limiting the impact of noise, it is impossible to extract any physical information from the detector's readout. While the field of gravitational-wave data analysis is vast and many different tasks can be achieved with different tools, in this dissertation we focus on the problem of searching a particular signal inside the detector's output. As the noise is typically louder than most of signals that we can conceivably detect, we have to employ sophisticated computational techniques to identify a signal inside the data stream. In other words, you cannot find a gravitational wave unless you are specifically looking for it.

One important search method goes under the name of matched filtering and consists in explicitly searching inside the data for a *particular* source, i.e. a particular binary. The search is done by computing the correlation between the data and the expected gravitational-wave signal emitted by the source. Of course, there are many possible binaries that could have emitted a detectable signal and, for this reason, we want to make sure that we don't miss any source. Therefore, we typically simultaneously search for millions of different sources inside the data. A little bit of jargon: we call *template* a signal that we search in the data produced by a given source (i.e. with given masses and spins), while the set of signal templates that are searched in the data is called *template bank*. Whenever a template shows a large correlation

value with the data, a search pipeline launches a *trigger*. In summary, to perform a matched-filter search, an interested experiment must (i) generate a template bank and (ii) wait until a computer filters the interferometer data with all the templates of the template bank.

Matched-filtering searches have proved their value by detecting more than 90 GW signals from the coalescence of compact objects but they are almost blind to signals that are not explicitly included inside the template bank, i.e. inside the set of signals to search in the data. This brings us to the problem addressed by this dissertation. Traditional searches typically search for signals with aligned spins, which are signals not showing precession. This clearly poses a limitation on their ability to detect precessing binaries and indeed, aligned-spin searches brought to the detection of very few precessing signals. Therefore we ask ourselves: is the observed lack of strongly precessing signals due to their rarity? Or, rather, is it caused by the limited sensitivity of current searches? To answer these questions and to make sure we are not currently missing any precessing signal in the data, we developed a matched-filtering search, based on the GstLAL pipeline, targeted to precessing signals and, for the first time, we searched real data looking for gravitational waves emitted by heavily precessing systems.

Our contributions: novel template placement method Despite being in principle straightforward, developing a search for precessing signals must face a number of technical challenges, whose solution makes up the original content of the dissertation. A first challenge comes with the generation of a template bank of precessing signals. Traditional template placement methods iteratively propose a new signal as a possible template. The proposal is then added to the bank only if its “distance” (defined in some meaningful way) with the other templates is larger than a certain threshold. Now, the problem with this method is that the number of distance computations to be performed increases with the number of parameters that characterize a template—the dimensionality of the space. An aligned-spin template is characterized by only four parameters (the two BH masses and the two spin components aligned to the orbital plane). On the other hand, a precessing parameter can be characterized by up to ten parameters and the traditional method becomes unfeasibly expensive in the precessing scenario.

To solve this challenge, we drop the expensive distance computation and instead we randomly draw template from a suitable, theoretically motivated distribution—this is not as bad as it might seem: a theorem guarantees that in a high number of dimension this approach is close to optimality. The sampling from the distribution is made feasible thanks to a machine learning model, designed for the purpose. While

our method delivers template banks up to 50% larger than those generated by the traditional method, it enables very fast bank generation (hours vs weeks). The speed of our method allows us to quickly identify a convenient region of the precessing parameter space, which is interesting from a physical point of view and at the same time can be covered with a relatively low number of templates. The template banks generated are used to search for precessing signals in the real interferometer data.

Our contributions: novel signal consistency test A second technical challenge inherent to precessing searches involves a so-called *signal consistency test*. During the filtering procedure the pipeline may occasionally launch spurious triggers, caused by instrumental noise artifacts. For this reason, each pipeline has to deploy a test to distinguish spurious triggers from triggers of astrophysical origin. The current version of the pipeline deploys a version of the test which is unsuitable for precessing signals. For this reason, we are forced to update the signal consistency test for precessing systems and to validate its performance. The updated test has been incorporated into the precessing search we developed.

Our contributions: sensitivity improvement and O3 search Having solved the two technical challenges above, we are ready to deploy a matched-filtering search for precessing signals. To assess the improvement over an aligned-spin search, we identify two different regions of the parameter space, one “Low q ” composed by systems with a limited mass asymmetry (mass ratio up to 5/1) and another “High q ” targeting systems with a large mass asymmetry (up to 12/1). For each of these two regions, we conduct both an aligned-spin search and a precessing search and we measured their sensitivity (evaluated in terms of the spacetime volume surveyed). We found that the precessing search performed *worse* than the aligned-spin one in the “Low q ” region: this can be explained by considering the larger number of templates of the precessing search, which increases the number of false positive triggers, hence harming the overall sensitivity. The situation is reversed in the “High q ” region, where for system with largely mis-aligned spins a sensitivity improvement of up to 120% was recorded. This is the primary takeaway from this investigation: in order to achieve a sensitivity improvement when searching for precessing signals, it is necessary to consider very asymmetric systems with largely mis-aligned spins!

Leveraging our results, we searched for precessing signals the public data collected by the LIGO observatory during the third observing run O3. Our search did not report any new discovery, however it allowed to place an upper limit on the detection rate of a population of precessing signals, not detected by past searches. Since our rates are consistent with the findings of previous analysis, we are able to confirm that precessing signals are indeed “as rare” as previously thought.

Despite bringing an interesting null result, our work does not end here. Indeed, other corners of the parameter space are left to be explored and other physical effects are to be incorporated within the matched-filtering searches. Due to its large computational cost, our approach is likely not suitable to be scaled to a larger scale search. For this reason, it is important to develop the next generation of matched-filtering searches, able to incorporate a large variety of physical effects at a moderate costs. The lessons learnt in this dissertation will provide invaluable guidance towards this goal.

Openbare Samenvatting

Algemene relativiteitstheorie, zwarte gaten en zwaartekrachtsgolven De theorie van de algemene relativiteit daagt ons conventionele begrip van concepten zoals tijd, energie en ruimte uit en wordt beschouwd als een nauwkeurige beschrijving van fysische verschijnselen die massieve objecten in het universum omvatten, zoals planeten of sterren. Achter de ingewikkelde wiskunde voorspelt de theorie het bestaan van een klasse van zeer dichte, sterachtige objecten, genaamd zwarte gaten (BH), die niet worden voorspeld door de klassieke Newtoniaanse zwaartekrachttheorie. Een zwart gat is een regio in de ruimte waar de zwaartekracht zo sterk is dat niets genoeg snelheid kan hebben om aan de zwaartekracht te ontsnappen. Daarom zal alles wat een zwart gat binnengaat, inclusief licht, voor altijd binnen een zwart gat blijven².

Hoewel zwarte gaten de bekendste voorspelling van de theorie zijn, zijn ze zeker niet de enige. De theorie voorspelt ook het bestaan van zwaartekrachtsgolven, die voortkomen als een natuurlijke consequentie van de eindige voortplantingssnelheid van de zwaartekrachtinteractie. Een simpel gedachtenexperiment kan dit verduidelijken. Stel je voor dat je een grote massa in een lege ruimte (in vacuüm) plaatst en met een kleine testmassa de zwaartekrachtseffecten ervan meet, die volgens schoolboeken worden waargenomen als een aantrekkingskracht naar de grote massa. Stel je nu voor dat (om welke reden dan ook) de massa op tijdstip t_0 begint te bewegen: wat gebeurt er dan met de testmassa? Of beter gezegd, wanneer voelt deze de verandering in de aantrekkingskracht vanwege de verplaatsing van de massa? Volgens de klassieke zwaartekrachttheorie is de interactie onmiddellijk: zodra de grote massa beweegt, voelt de kleine massa het effect. Natuurlijk weten we dat dit beeld naïef is, aangezien niets kan reizen met “oneindige” snelheid. Om deze reden moet de algemene relativiteit voorspellen dat de testmassa een variatie in de zwaartekracht pas op een vertraagd tijdstip t_R waarneemt. Wat gebeurt er tussen t_0 en t_R ? De zwaartekrachtinteractie plant zich voort in de vorm van een zwaartekrachtsgolf! Dit beeld is

²Dit motiveert hun naam: aangezien er geen licht wordt uitgezonden of gereflecteerd, verschijnen ze voor elke waarnemer die naar hen kijkt als een puur zwarte bol.

natuurlijk sterk vereenvoudigd en er zijn veel nuances over wat je wel of niet kunt definiëren als een zwaartekrachtsgolf. Desondanks blijft het belangrijkste fysische idee overeind: een zwaartekrachtsgolf is een zich voortplantende zwaartekrachtinteractie, net zoals licht een zich voortplantende elektromagnetische interactie is.

Zwarte gaten-binaire systemen Wanneer we vaststellen dat zwaartekrachtsgolven bestaan (merk op dat deze stap niet triviaal is, aangezien het bijna vijftig jaar van wetenschappelijk debat kostte om de gemeenschap hiervan te overtuigen), kunnen we ons afvragen: hoe worden zwaartekrachtsgolven geproduceerd? Duiken in alle wiskundige details van de theorie biedt een vrij eenvoudig antwoord: zwaartekrachtsgolven worden geproduceerd door versnellende objecten. Natuurlijk zijn er veel versnellende objecten in het universum, waaronder een trein die een station verlaat, maar een bijzonder intrigerende klasse van bronnen van zwaartekrachtsgolven komt van binaire systemen van compacte objecten. Hier zijn enkele definities op hun plaats: een compact object is een zeer dicht, sterachtig object, zoals een neutronenster of een zwart gat, terwijl een compact binair systeem een paar compacte objecten is die om elkaar heen draaien onder invloed van hun wederzijdse zwaartekracht. Simpel gezegd is een compact binair systeem vergelijkbaar met het bekende systeem van de aarde en de zon, maar dan vervangen door zwaardere en dichtere objecten. We noemen een systeem van twee zwarte gaten een binair zwart gat. Beide objecten bevinden zich in een (bijna) uniforme cirkelvormige beweging en, nog belangrijker, ze ervaren een centripetale versnelling: dit betekent dat ze zwaartekrachtsgolven uitzenden!

Vanwege hun belang als bronnen van zwaartekrachtsgolven wordt de fysica van **een binair zwart gat** uitgebreid bestudeerd in de meeste van hun details en zijn er inderdaad zeer nauwkeurige voorspellingen voor het GW-signalen dat door het systeem wordt uitgezonden. Ondanks de wiskundige complexiteit ontstaan enkele eenvoudige fysische ideeën. Allereerst is het algemeen bekend dat, aangezien de zwaartekrachtraling energie transporteert, de twee objecten de energieverliezen compenseren door steeds dichter naar elkaar te draaien, totdat ze uiteindelijk samensmelten tot één enkel object. Bovendien spelen de spins (d.w.z. de rotatie rond een bepaalde as) van de twee objecten een rol in de dynamiek van het systeem, waardoor het uitzenden van zwaartekrachtsgolven wordt vertraagd of versneld. Ten slotte, als de twee spins niet zijn uitgelijnd met elkaar of met het baanvlak, ondergaat het systeem precessie. Dit betekent dat het baanvlak zelf wiebelt en van oriëntatie verandert terwijl de twee zwarte gaten naar elkaar toe spiralen, wat unieke kenmerken in de zwaartekrachtsgolf-signalen introduceert.

De observatie van de zwaartekrachtsgolven uitgezonden door de samensmelting van compacte objecten stelt wetenschappers in staat om een grote verscheidenheid

aan belangrijke wetenschappelijke vragen te beantwoorden. Allereerst kan men zich afvragen hoeveel binaire systemen er in het universum bestaan, wat hun eigenschappen zijn en hoe ze zijn ontstaan—dit valt onder het domein van zogenaamde populatiestudies. Een andere set vragen betreft de geldigheid van de theorie van de algemene relativiteit om dergelijke ingewikkelde binaire systemen te beschrijven—dit is het domein van de testen van de algemene relativiteitstheorie. Bovendien is het mogelijk om de eigenschappen van dichte materie te beperken door de samensmelting van neutronensterren en de complexe fenomenologie die daarmee gepaard gaat te observeren. Ten slotte, omdat de bronnen zich op zeer grote afstanden bevinden, kunnen we ze gebruiken om de structuur van ons universum te bestuderen, waardoor we beperkingen kunnen opleggen aan de expansiesnelheid en de geschiedenis ervan, en zo bijdragen aan het vakgebied van de kosmologie.

Detecteren van zwaartekrachtsgolven van zwarte gaten binaries De bovenstaande lijst geeft slechts een hint van het wetenschappelijke potentieel van het observeren van compacte binairen, en het is dan ook geen verrassing dat er miljarden euro's zijn toegewezen om grote instrumenten te bouwen die zijn gewijd aan de detectie van zwaartekrachtsgolven, zoals LIGO, Virgo of KAGRA. Dergelijke instrumenten zijn ontworpen om een kleine variatie in de afstand tussen twee testmassa's te meten, wat het effect is van een zwaartekrachtsgolf op de ruimtetijd zoals voorspeld door de algemene relativiteitstheorie. Een dergelijke nauwkeurige meting kan alleen worden bereikt met zeer grote interferometers, door het interferentiepatroon te analyseren van twee lasers die worden afgevuurd in de kilometerslange armen. De interferometer levert een tijdsafhankelijke schatting van de differentiële armlengte (d.w.z. het verschil tussen de lengte van de twee armen), die nauw verwant is aan elke inkomende zwaartekrachtsgolf.

Met de theorie van compacte binaire systemen in handen en met geavanceerde instrumenten die werken om de zwaartekrachtsgolven te meten, zou het lijken alsof het moeilijkste deel van het werk is gedaan. Helaas is het grootste deel van de detectoroutput ruis. Zonder de juiste data-analysetechnieken om de impact van ruis te beperken, is het onmogelijk om enige fysieke informatie uit de detectorlezing te halen. Hoewel het veld van zwaartekrachtsgolfsdata-analyse breed is en veel verschillende taken kan worden uitgevoerd met verschillende tools, richten we ons in deze dissertatie op het probleem van het zoeken naar een bepaald signaal in de detectoroutput. Aangezien de ruis sterker is dan elk signaal dat we kunnen detecteren, moeten we geavanceerde computatietechnieken gebruiken om een signaal in de datastream te identificeren. Met andere woorden, je kunt een zwaartekrachtsgolf niet vinden tenzij je er specifiek naar op zoek bent.

Een belangrijke zoekmethode staat bekend als matched filtering en bestaat uit het expliciet zoeken in de gegevens naar een specifieke bron, d.w.z. een specifiek binair systeem. De zoekopdracht wordt uitgevoerd door de correlatie te berekenen tussen de gegevens en het verwachte zwaartekrachtsgolfsignaal dat door de bron wordt uitgezonden. Natuurlijk zijn er veel mogelijke binairen die een detecteerbaar signaal hadden kunnen uitzenden en daarom willen we ervoor zorgen dat we geen enkele bron missen. Daarom zoeken we doorgaans tegelijkertijd naar miljoenen verschillende bronnen in de gegevens. Een beetje vaktaal: we noemen een signaal dat we in de gegevens zoeken, geproduceerd door een bepaalde bron (d.w.z. met bepaalde massa's en spins), een *sjabloon*, terwijl de set signaalsjablonen die in de gegevens worden doorzocht, de *sjabloonbank* wordt genoemd. Telkens wanneer een sjabloon een hoge correlatielijnige waarde met de gegevens vertoont, start een zoekpipeline een *trigger*. Samengevat, om een matched-filter-zoekopdracht uit te voeren, moet een geïnteresseerd experiment i) een sjabloonbank genereren en ii) wachten tot een computer de interferometergegevens met alle sjablonen van de sjabloonbank heeft gefilterd.

Matched-filtering zoekopdrachten hebben hun waarde bewezen door meer dan 90 GW-signalen te detecteren van de samensmelting van compacte objecten, maar ze zijn bijna blind voor signalen die niet expliciet zijn opgenomen in de sjabloonbank, d.w.z. binnen de reeks signalen om in de gegevens te zoeken. Dit brengt ons bij het probleem dat in deze dissertatie wordt behandeld. Traditionele zoekopdrachten zoeken doorgaans naar signalen met uitgelijnde spins, wat signalen zijn die geen precessie vertonen. Dit vormt duidelijk een beperking op hun effectiviteit bij het detecteren van precesserende binairen, en inderdaad hebben zoekopdrachten met uitgelijnde spins geleid tot de detectie van zeer weinig precesserende signalen. Daarom stellen we ons de vraag: is het waargenomen gebrek aan sterk precesserende signalen te wijten aan hun zeldzaamheid? Of wordt het eerder veroorzaakt door de beperkte gevoeligheid van de huidige zoekopdrachten? Om deze vragen te beantwoorden en ervoor te zorgen dat we momenteel geen precesserende signalen in de gegevens missen, hebben we een matched-filter-zoekopdracht ontwikkeld, gebaseerd op de GstLAL-pipeline, gericht op precesserende signalen en, voor het eerst, hebben we echte gegevens doorzocht op zoek naar zwaartekrachtsgolven die worden uitgezonden door sterk precesserende systemen.

Onze bijdragen: nieuwe methode voor sjabloonplaatsing Hoewel het in principe eenvoudig lijkt, brengt het ontwikkelen van een zoekmethode voor precesserende signalen een aantal technische uitdagingen met zich mee, waarvan de oplossing de originele inhoud van het proefschrift vormt. Een eerste uitdaging is de generatie van een sjabloonbank voor precesserende signalen. Traditionele methoden voor sjabloon-

plaatsing stellen iteratief een nieuw signaal voor als mogelijke sjabloon. Dit voorstel wordt alleen aan de bank toegevoegd als de “afstand” (op een zinvolle manier gedefinieerd) tot de andere sjablonen groter is dan een bepaalde drempelwaarde. Het probleem met deze methode is dat het aantal berekeningen van afstanden toeneemt met het aantal parameters dat een sjabloon kenmerkt—de dimensionaliteit van de ruimte. Een sjabloon met uitgelijnde spin wordt gekarakteriseerd door slechts vier parameters (de twee massa’s van de zwarte gaten en de twee spincomponenten die zijn uitgelijnd met het baanvlak). Aan de andere kant kan een precesserend signaal worden gekarakteriseerd door maximaal tien parameters, waardoor de traditionele methode onhaalbaar duur wordt in het geval van precessie.

Om deze uitdaging op te lossen, laten we de dure afstandsberekening vallen en trekken we in plaats daarvan willekeurig sjablonen uit een geschikte, theoretisch onderbouwde verdeling—dit is niet zo nadelig als het lijkt: een stelling garandeert dat deze aanpak in een groot aantal dimensies dicht bij optimaliteit ligt. Het bemonsteren van de verdeling wordt mogelijk gemaakt door een machine learning-model, speciaal hiervoor ontworpen. Hoewel onze methode sjabloonbanken oplevert die tot 50% groter zijn dan die welke worden gegenereerd door de traditionele methode, maakt het zeer snelle bankgeneratie mogelijk (uren versus weken). De snelheid van onze methode stelt ons in staat om snel een geschikt gebied van de precessieparameterruimte te identificeren, dat interessant is vanuit fysiek oogpunt en tegelijkertijd met een relatief klein aantal sjablonen kan worden afgedekt. De gegenereerde sjabloonbanken worden gebruikt om precesserende signalen in echte interferometergegevens te zoeken.

Onze bijdragen: nieuwe consistentietest voor signalen Een tweede technische uitdaging bij precessie-zoektochten betreft de zogenoemde *consistentietest voor signalen*. Tijdens de filterprocedure kan de pijplijn soms valse triggers activeren, veroorzaakt door ruisartefacten van instrumenten. Om deze reden moet elk analyse-algoritme een test inzetten om valse triggers te onderscheiden van triggers van astrofysische oorsprong. De huidige versie van het algoritme maakt gebruik van een versie van de test die ongeschikt is voor precesserende signalen. Daarom zijn we genoodzaakt de consistentietest voor precessiesystemen te updaten en de prestaties ervan te valideren. De bijgewerkte test is geïntegreerd in de precessie-zoektocht die we hebben ontwikkeld.

Onze bijdragen: verbetering van de gevoeligheid en O3-zoektocht Nadat we de twee hierboven genoemde technische uitdagingen hebben opgelost, zijn we klaar om een matched-filtering zoekopdracht naar precessesignalen uit te voeren. Om de verbetering ten opzichte van een zoektocht met uitgelijnde spins te beoordelen, identifi-

ceren we twee verschillende gebieden van de parameterspace: een “Laag q ” dat bestaat uit systemen met een beperkte massa-asymmetrie (massaverhouding tot 5/1) en een andere “Hoog q ” die zich richt op systemen met een grote massa-asymmetrie (tot 12/1). Voor elk van deze twee gebieden voeren we zowel een uitgelijnde-spins zoekopdracht als een precessiezoekopdracht uit en meten we hun gevoeligheid (geëvalueerd in termen van het gesurveilleerde ruimtetijdvolume). We ontdekten dat de precessiezoekopdracht *slechter* presteerde dan de uitgelijnde-spin zoektocht in het “Laag q ”-gebied: dit kan worden verklaard door het grotere aantal sjablonen van de precessiezoekopdracht, wat het aantal vals-positieve triggers verhoogt en dus de algehele gevoeligheid schaadt. De situatie is omgekeerd in het “Hoog q ”-gebied, waar voor systemen met sterk niet-uitgelijnde spins een verbetering van de gevoeligheid tot 120% werd geregistreerd. Dit is de belangrijkste conclusie van dit onderzoek: om een verbetering van de gevoeligheid te bereiken bij het zoeken naar precessiesignalen, is het noodzakelijk om zeer asymmetrische systemen met sterk niet-uitgelijnde spins in overweging te nemen!

Met behulp van onze resultaten hebben we gezocht naar precessiesignalen in de openbare gegevens die door het LIGO-observatorium zijn verzameld tijdens de derde waarnemingsronde O3. Onze zoektocht heeft geen nieuwe ontdekking gerapporteerd, maar het stelde ons in staat om een bovengrens vast te stellen voor het detectietempo van een populatie van precessiesignalen die niet door eerdere zoektochten zijn gedetecteerd. Aangezien onze tarieven consistent zijn met de bevindingen van eerdere analyses, kunnen we bevestigen dat precessiesignalen inderdaad “zo zeldzaam” zijn als eerder werd gedacht.

Ondanks het brengen van een interessant nulresultaat, eindigt ons werk hier niet. Inderdaad, andere uithoeken van de parameterruimte moeten nog worden verkend en andere fysische effecten moeten worden opgenomen binnen de matched-filtering zoekopdrachten. Vanwege de grote computationele kosten is onze benadering waarschijnlijk niet geschikt om op grotere schaal te worden opgeschaald. Om deze reden is het belangrijk om de volgende generatie matched-filtering zoekopdrachten te ontwikkelen, die in staat zijn om een grote verscheidenheid aan fysische effecten op gematigde kosten te integreren. De lessen die in deze dissertatie zijn geleerd, zullen van onschabare waarde zijn voor deze doelstelling.

Selected Publications

Precessing Searches

- [1] Stefano Schmidt, Bhooshan Gadre, and Sarah Caudill. Gravitational-wave template banks for novel compact binaries. *Phys. Rev. D*, 109(4):042005, 2024
- [2] Stefano Schmidt and Sarah Caudill. Novel signal-consistency test for gravitational-wave searches of generic black hole binaries. *Phys. Rev. D*, 110(2):023042, 2024
- [3] Stefano Schmidt et al. Searching for gravitational-wave signals from precessing black hole binaries with the GstLAL pipeline. *Phys. Rev. D*, 110(2):023038, 2024
- [4] Stefano Schmidt et al. Searching for Asymmetric and Heavily Precessing Binary Black Holes in the Gravitational Wave Data from the LIGO Third Observing Run. *Phys. Rev. Lett.*, 133(20):201401, 2024

Machine Learning Waveform Modeling

- [5] Stefano Schmidt, Matteo Breschi, Rossella Gamba, Giulia Pagano, Piero Retegno, Gunnar Riemenschneider, Sebastiano Bernuzzi, Alessandro Nagar, and Walter Del Pozzo. Machine Learning Gravitational Waves from Binary Black Hole Mergers. *Phys. Rev. D*, 103(4):043020, 2021
- [6] Jacopo Tissino, Gregorio Carullo, Matteo Breschi, Rossella Gamba, Stefano Schmidt, and Sebastiano Bernuzzi. Combining effective-one-body accuracy and reduced-order-quadrature speed for binary neutron star merger parameter estimation with machine learning. *Phys. Rev. D*, 107(8):084037, 2023
- [7] Tim Grinbergen, Stefano Schmidt, Chinmay Kalaghatgi, and Chris van den Broeck. Generating higher order modes from binary black hole mergers with machine learning. *Phys. Rev. D*, 109(10):104065, 2024

Miscellaneous

- [8] Paloma Laguarta et al. Detection of anomalies amongst LIGO's glitch populations with autoencoders. *Class. Quant. Grav.*, 41(5):055004, 2024
- [9] Alessandro Martini, Stefano Schmidt, Gregory Ashton, and Walter Del Pozzo. Maximum entropy spectral analysis: an application to gravitational waves data analysis. *Eur. Phys. J. C*, 84(10):1023, 2024

List of Figures

1.1	Effect on a ring of test masses of a monochromatic gravitational wave, propagating perpendicularly to the plane of the particles. For different phases ωt of the wave, we report the position of the particles, for the a plus-polarized waveform (top line) and a cross-polarized waveform (bottom line). Each particle moves according to Eq. (1.37) and Eq. (1.38) under the effect of the plus and cross polarization respectively.	10
1.2	Waveform emitted by a binary system according to the Newtonian approximation Eqs. (1.85-1.86) (in blue), compared to a more realistic prediction generated with the state-of-the-art <code>IMRPhenomT</code> waveform approximant (underlaid in orange). The two waveforms refer to a system with masses $m_1 = 30 M_\odot$ and $m_2 = 10 M_\odot$, seen from a distance of 100 Mpc and with zero inclination, and generated from a minimum frequency $f_{\min} = 5$ Hz.	20
1.3	Schematic representation of the orbital plane of a precessing binary system with the orbital angular momentum \mathbf{L} , the total spin vector \mathbf{S} and the total angular momentum \mathbf{J} . We also report the angles α and β , which parameterize the precession of the orbital angular momentum \mathbf{L} around the \mathbf{J} -axis with respect to an inertial Cartesian frame.	24
1.4	The Euler angles α and β obtained by numerically integrating Eqs. (1.96-1.97) (in blue) and with their approximate expressions Eq. (1.89) and Eq. (1.88). The fast oscillations around the approximate expression for β are caused by the nutation of \mathbf{L} along its precession motion. The angles refer to a system with masses $m_1 = 15 M_\odot$ and $m_2 = 1.5 M_\odot$ and with dimensionless spins $s_1 = (0.8, 0.0, -0.5)$ and $s_2 = (-0.7, 0.1, 0.3)$.	26

4.1	Study of the accuracy for several normalizing flow, trained on different manifolds. For each manifold, we compute the logarithmic ratio $\log_{10} \frac{p^{\text{flow}}}{p^{\text{true}}}$ between the PDF computed by the flow and the true one. We use 40000 test points from the validation set of each manifold. Details on the manifold considered are reported in Tab. 4.1.	88
4.2	Validation of the random template placement algorithm. For three of the manifolds introduced in Tab. 4.1, we plot the number the number of templates $N_{\text{templates}}$ of a random template bank as a function of the number of live points $N_{\text{livepoints}}$ used to estimate the covering fraction. For each template bank, we set $\eta = 0.9$ and $MM = 0.97$	88
4.3	Fitting factor studies for the two template banks introduced in Sec. 4.4. As discussed in Sec. 4.4.1 and 4.4.2 respectively, “HM bank” is designed to reproduce [328] and targets high mass non-spinning systems with HM content, while the “All-sky bank” bank covers aligned-spin systems (without HM) over a broad mass range, following [380]. We report the cumulative histogram of the fitting factors of 10^5 injections samples across the parameter space.	90
4.4	Validation of the two template banks designed to reproduce two independently generated banks. On the left side, we report results pertaining the “HM bank”, while on the right side we report the “All-sky bank”. For each two dimensional bin, we report the median fitting factor of 10^5 injections covering the parameter space, as described in the text.	90
4.5	Corner plot with the templates of the precessing bank described in Sec. 4.5.1. Along the diagonals, we show the histogram of the template number as a function of each coordinate.	94
4.6	Accuracy of the normalizing flow trained used to generate the precessing bank in Sec. 4.5.1. The accuracy is expressed in terms of the logarithmic ratio between the template density PDF p^{true} Eq. (4.10) and its approximation p^{flow} given by the flow. The flow accuracy is evaluated on 40000 test points.	95
4.7	Fitting factor study of the precessing bank, introduced in Sec. 4.5.1. For each bin, we color-code the median fitting factor of 10^5 injections sampled “On manifold”, as described in the text.	95

4.8	Fitting factor study of the precessing bank of Sec. 4.5.1. Unlike Fig. 4.7, here we focus on the low q , low M region, where the random placement method fails. For each bin, we color-code the median fitting factor of 5×10^4 injections sampled “On manifold”, as described in the text.	96
4.9	Cumulative fitting factor for the precessing bank introduced in Sec. 4.5.1. The 10^5 injections “Full precessing” have isotropic spins, while the 3×10^5 precessing injections “On manifold” are sampled on the manifold <code>logMq_s1xz_s2z_iota</code> and they have $s_{1y} = s_{2x} = s_{2y} = \phi = 0$. For the injections “On manifold”, we plot separately the low q , low M corner, characterized by $q \leq 1.2$ and $M \leq 20 M_\odot$. The other two histograms exclude this region.	98
4.10	Determinant of the metric $ M $ as a function of mass ratio q for different values of s_1 . The metric is evaluated on the manifold <code>Mq_s1xz_s2z_iota</code> , with $M = 10 M_\odot$, $\theta_1 = \pi/2$, $s_{2z} = -0.3$ and $\iota = \pi/2$. It is manifest that in all cases, the metric determinant vanishes while $q \rightarrow 1$	99
4.11	Accuracy of the normalizing flow trained used to generate the aligned-spin HM bank of Sec. 4.5.2. The accuracy is expressed in terms of the logarithmic ratio between the template density PDF p^{true} Eq. (4.10) and its approximation p^{flow} given by the flow. The flow accuracy is evaluated on 40000 test points.	100
4.12	Fitting factor study of the aligned-spin HM bank, introduced in Sec. 4.5.2. For each bin, we color-code the median fitting factor of 10^5 injections sampled uniformly from the parameter space.	101
4.13	Corner plot with the templates of the aligned-spin HM bank described in Sec. 4.5.2. Along the diagonals, we show the histogram of the template number as a function of each coordinate.	102
4.14	Cumulative fitting factor for the aligned-spin HM bank described in Sec. 4.5.2. The histogram is built upon 10^5 injections sampled from the manifold.	103

4.15 Study of the size of a template bank in the neutron star-black hole parameter space. Each point refers to a template bank on the manifold `logMq_s1xz_iota`, covering a total mass range $M \in [M_{\min}, M_{\max}]$. The component masses are limited to $m_1 \in [10, 60] M_{\odot}$ and $m_2 \in [1, 3] M_{\odot}$, with mass ratios $q \in [3.3, 15]$. In the plot we report the number of templates $N_{\text{templates}}$ as a function of the maximum total mass M_{\max} , for different minimal match requirements. The resulting banks are huge, with tens of millions of templates, showing that a search for precessing NSBH binaries is still prohibitively costly. 104

4.16 For each eigenvector of the metric, we compute the empirical relation between the mis-match $1 - \mathcal{M}$ and the distance ϵ of points along the eigenvector direction. The solid line shows the relation predicted by the metric, while the dashed line shows a parabolic fit. In the legend are reported the quadratic coefficients of both lines. 112

5.1 Predicted and measured absolute value of the SNR timeseries, $|R|$ and ρ respectively, for a precessing signal in zero noise and with unit magnitude, filtered with a perfectly matching template. We measure the SNR using both the “standard” SNR and the “symphony” SNR. The predicted SNR is computed with three different prescriptions, R_{std} Eq. (5.8) suitable for the “standard” SNR, and R_{sym} Eq. (5.18) and R_{mix} Eq. (5.24), both suitable for the “symphony” SNR. We also report in the bottom panel the difference between the expected and measured SNR timeseries. In the bottom panel we report the difference between both $|R_{\text{sym}}|$ or $|R_{\text{mix}}|$ and the “symphony” SNR. It is manifest that $|R_{\text{sym}}|$ and ρ_{sym} show perfect agreement between each other while $|R_{\text{std}}|$ and $|R_{\text{mix}}|$ do not accurately predict the relevant SNR timeseries. Also note that the peak of “standard” search statistic is lower than unity, meaning that performing matched filtering with templates Eq. (3.21) is not able to fully recover the SNR of a precessing signal. The signal is injected into Gaussian noise, sampled from the LIGO Livingston PSD [439] with a rate of 4096 Hz. The waveform is characterized by masses $m_1, m_2 = 28 M_{\odot}, 3 M_{\odot}$ and spins $\mathbf{s}_1 = (-0.8, 0.02, -0.5)$ and $\mathbf{s}_2 = 0$, observed with an inclination $\iota = 2.66$. It was generated starting from a frequency of 12 Hz with the approximant `IMRPhenomXP` [47]. 121

5.2 Values of ξ^2 Eq. (5.4) as a function of the absolute value of $(\hat{h}_+ \hat{h}_\times)$, which quantifies the precession and/or HM content of a signal. Each value is computed on random <i>precessing</i> BBHs, injected into Gaussian noise at a constant SNR. The left and center panels show SNRs of 10 and 100 respectively, while the right panel tackles the case of zero noise with injected signal normalized to 1. ξ^2 is computed using three different prescriptions. ξ_{std}^2 is obtained from $R_{\text{std}}(t)$ and $z_{\text{std}}(t)$ (label “standard”). ξ_{sym}^2 is computed using $R_{\text{sym}}(t)$ and $z_{\text{sym}}(t)$ (label “symphony”), while ξ_{mix}^2 uses $R_{\text{mix}}(t)$ and $z_{\text{sym}}(t)$. For this study we set ACL = 701.	123
5.3 Values of ξ^2 Eq. (5.4) as a function of the absolute value of $(\hat{h}_+ \hat{h}_\times)$, which quantifies the precession and/or HM content of a signal. Each value is computed on random <i>aligned-spin</i> BBHs with HM content, injected into Gaussian noise at a constant SNR. The left and center panels show SNRs of 10 and 100 respectively, while the right panel tackles the case of zero noise with injected signal normalized to 1. ξ^2 is computed using three different prescriptions. ξ_{std}^2 is obtained from $R_{\text{std}}(t)$ and $z_{\text{std}}(t)$ (label “standard”). ξ_{sym}^2 is computed using $R_{\text{sym}}(t)$ and $z_{\text{sym}}(t)$ (label “symphony”), while ξ_{mix}^2 uses $R_{\text{mix}}(t)$ and $z_{\text{sym}}(t)$. For this study we set ACL = 701.	125
5.4 Performance of the “standard” test as a function of the parameter space. We color each bin according to the median value of ξ_{std}^2 and we consider the mass ratio q , the polar spin components s_1, θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text.	126
5.5 Comparison between the “mixed” and the “standard” signal consistency tests as a function of the parameter space. We color each bin according to the median value of $\xi_{\text{mix}}^2/\xi_{\text{std}}^2$ and we consider the mass ratio q , the polar spin components s_1, θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text.	128
5.6 Comparison between the “mixed” and the “standard” signal consistency tests as a function of the parameter space. We color each bin according to the median value of $\xi_{\text{mix}}^2/\xi_{\text{sym}}^2$ and we consider the mass ratio q , the polar spin components s_1, θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text.	128

5.7	Values of ρ^A as a function of the parameter space. We color each bin according to the median value of ρ^A and we consider the mass ratio q , the polar spin components s_1 , θ_1 and of the inclination angle ι of the 15000 test BBHs described in the text. ρ^A is intended as a measure of the goodness of the “mixed” signal consistency test and as such, it correlate very well with the values of ξ_{mix}^2 (see also Fig. 5.6).	128
5.8	Numerical stability of the approximant <code>IMRPhenomXP</code> in the aligned-spin limit as a function of the inclination angle ι . We compute the values of ρ^A and for a BBH with total mass $M = 10 M_\odot$ and mass ratio $q = 8$ and with spins $s_{1z}, s_{2z} = -0.4, 0.6$, seen at different inclination angles ι . We repeat the experiment with both the precessing approximant <code>IMRPhenomXP</code> and the aligned-spin approximant <code>IMRPhenomXAS</code> . While for the aligned-spin approximant ρ^A are both zero up to numerical precision, for the precessing approximant the two quantities become non zero as the inclination gets close to $\iota \simeq \pi/2$. In this case, since $h_x \rightarrow 0$, the numerical noise introduced by the spin twist procedure affects the normalized cross polarization \hat{h}_x	130
5.9	Discrepancies between the “mixed” signal consistency test ξ_{mix}^2 Eq. (5.24) and a possible alternative definition Eq. (5.27) $\xi_{\text{mix-bis}}^2$. For the 15000 BBH signals introduced in the text, we report a histogram with the fractional difference $\frac{\xi_{\text{mix-bis}}^2 - \xi_{\text{mix}}^2}{\xi_{\text{mix}}^2}$ between the two consistency tests. The discrepancies between the two versions are negligible, showing that Eq. (5.24) and Eq. (5.27) are mostly equivalent.	130
5.10	Values of ξ^2 for different window lengths (ACL) used for the integral in Eq. (5.4). For each ACL, we report the values of ξ_{std}^2 (blue), ξ_{mix}^2 (orange) and ξ_{sym}^2 (green), following the color code introduced in Figs. 5.2 and 5.3. The data refers to 15000 <i>precessing</i> signals injected at $\text{SNR} = 100$	131
5.11	Performance in real detector’s noise of the three signal consistency tests discussed in the paper. We randomly select a number of precessing templates as in Fig. 5.2 and we inject a corresponding signal into both Gaussian noise and real detector’s noise. The detector’s noise was recorded by the LIGO-Hanford observatory between GPS times 1245708288 s and 1246756864 s. In each histogram, we report the values of ξ^2 for the Gaussian noise (orange) and Real noise (blue) case. Each panel refers to a different type of signal consistency test and to different values of the injected SNR.	132

6.1 Binned templates of the “Low q ” precessing template bank. For each equal size bin, we color code the logarithmic number of templates so that the color is also a measure of the template density. In the different panels, we consider the variables $M, q, s_1, \theta_1, \iota$. The bank was first introduced in [1]. 144

6.2 Binned templates of the “High q ” precessing template bank. For each equal size bin, we color code the logarithmic number of templates so that the color is also a measure of the template density. In the different panels, we consider the variables $M, q, s_1, \theta_1, \iota$ 144

6.3 Fitting factor study for the four template banks considered in this work. For each given value of the match, we report the fraction of injections with fitting factor lower than than that value. The two aligned-spin banks (AS) and the two precessing banks (P) are tested against the relevant injection set of fully precessing injections used for searching the data, with spins magnitudes reaching values up to 0.99. To evaluate their ability to cover the non-precessing space, the two AS template banks are also tested against a set of aligned-spin injections. The composition of the injection sets is described in the text. In the left panel, we report the results concerning the “Low q ” parameter space, while the right refers to the “High q ” parameter space. The vertical dashed lines mark the target minimal match of the template banks: 0.9 for the “High q ” precessing bank and 0.97 for the others. 145

6.4 Fitting factor of the “High q ” precessing template bank, as a function of total mass M , mass ratio q , tilt θ_1 and magnitude s_1 of the first BH, inclination angle ι and reference phase ϕ . For each bin, we report the 5th percentile of the fitting factor distribution. Note that while the fitting factor study covers values of primary spin between 0 and 0.99, the template bank only gathers templates with $s_1 \in [0.5, 0.9]$ 146

6.5 Compression achieved by the Singular Value Decomposition of the templates of the two precessing template banks considered in this work. For each bin, the compression is measured by the ratio between the number of filters and the number of templates. The number of filters is obtained by averaging, across different time slices, the number of SVD basis. The average is weighted by the sampling rate of each time slice. For different choices of the sorting variables (α^1, α^2) used to construct the bank split, we report the histogram of the compression achieved for different bins. 149

6.6 Histogram with the fractional SNR gain of a precessing search over its aligned-spin counterpart. For the two cases “Low q ” and “High q ”, we report the discrepancy $\Delta\text{SNR} = \text{SNR}_{\text{precessing}} - \text{SNR}_{\text{aligned-spin}}$ of the SNR recovered by the precessing and the aligned-spin, normalized by the SNR of each injection. The injections are described in Sec. 6.3 and the data refer to the LIGO Livingston detector. 151

6.7 Histogram with the difference between the ξ^2 measured by a precessing search and its aligned-spin counterpart. For the two cases “Low q ” and “High q ”, we report the discrepancy $\Delta\xi^2 = \xi^2_{\text{precessing}} - \xi^2_{\text{aligned-spin}}$, evaluated on the set of precessing injections described in Sec. 6.3. Data refer to the LIGO Livingston detector. 151

6.8 “SNR- ξ^2 ” plot for the “High q ” aligned-spin search. On the left panel, we report the SNR and ξ^2 of the background triggers, not associated to any injection, and the injection triggers, corresponding to an injection. On the right panel, we only report the injection triggers, colored by the logarithm of the FAR assigned by the pipeline. Data refer to the LIGO Livingston detector. The red star refers to the trigger produced by the GW event GW191204_171526 [18], detected by our search with $FAR < \frac{1}{1000} \text{ yr}^{-1}$ 152

6.9 “SNR- ξ^2 ” plot the “High q ” precessing search. On the left panel, we report the SNR and ξ^2 of the background triggers, not associated to any injection, and the injection triggers, corresponding to an injection. On the right panel, we only report the injection triggers, colored by the logarithm of the FAR assigned by the pipeline. Data refer to the LIGO Livingston detector. The red star refers to the trigger produced by the GW event GW191204_171526 [18], detected by our search with $FAR < \frac{1}{1000} \text{ yr}^{-1}$ 153

6.10 Efficiency of the BBH searches as a function of the loudest SNR measured by the three detectors, also called the “Decisive SNR”. The efficiency is defined as the fraction of signal detected at constant $FAR = 1/10$ years. In each panel, we show the efficiency pertaining the aligned-spin (AS) and the precessing (P) searches, reported separately for $\chi_P > 0.5$ and $\chi_P < 0.5$. The left panel refers to the “Low q ” search, while the right panel to the “High q ” search. 154

6.11 Sensitive volume increase as a function of the FAR detection thresh- old. The increase in sensitive volume is computed in terms of the VT ratio, measured between the VT of a precessing search and the VT of an aligned-spin searches. The two blue lines refer to the “Low q ” re- gion, while the blue lines correspond to “High q ”. A dashed line refers to the VT computed only from injections with $\chi_P > 0.5$, while a solid line corresponds to $\chi_P < 0.5$	155
6.12 Sensitive volume increase for different bins in the $q-\chi_P$ space. The left panel refers to the “Low q ” region while the right panel to the “High q ” region. For each bin, we compute the VT ratio between a precessing spin search and the corresponding aligned-spin search. The VT ratio is encoded into the color scale and is computed at a $FAR = 1/10$ years.	156
7.1 Median masses of the GWTC candidates as computed with parameter estimation. The GWTC candidates detected by our search are denoted by large circles, colored according to their FAR. A black cross marks the GWTC events that our search missed. The blue dots represent the templates used in our search bank.	165
7.2 Efficiency of the search as a function of the decisive SNR, i.e. the max- imum injected SNR between the two detectors. The efficiency is mea- sured on four different chunks with GPS start time and length reported on the caption separated by a dash.	166
7.3 Search sensitivity computed over different regions of the parameter space. Each dot marks the center of a set of log-normally distributed injections and it is colored according to the sensitive spacetime volume measured using such injection set. The number superimposed denotes the ratio between the sensitive spacetime volume of our search and the one of an aligned-spins search targeting the same space.	167

List of Tables

4.1	Details of the manifold considered for the validation of the normalizing flow model in Fig. 4.1. For each manifold, we report the variables being sampled together with their ranges. We also list the frequency range considered, the waveform approximant used, the number of dimensions D of the manifold as well as the number of hidden features for each layer of the flow.	86
4.2	Details of the two banks available in the literature that we reproduce with our code. For each bank, we indicate the parameter space considered and the approximant used. We also compare the number of templates of the banks obtained with the different methods.	89
6.1	We summarize here the most important features of the four template banks considered in this work. We select two regions of the parameter space, labeled “Low q ” and “High q ” respectively. For each region, we generate two template banks, one only gathering aligned-spin signals (AS) and another one including precessing signals (P). The aligned-spin banks covers the variable M, q, χ_{eff} , while the precessing template bank cover the variables $M, q, s_1, \theta_1, s_{2z}, \iota$. For each template bank we report the range covered by each variable as well as the number of templates (size). The BBH variables not explicitly mentioned in the table are set to 0.	142
7.1	Confident events detected by our search with $p_{\text{astro}} > 0.5$. For each event, we report the event name, together with the SNR and FAR. For reference, we also report the SNR _{GWTC} reported in the transient catalog. 164	

7.2 Candidates not reported in GWTC, detected by our search with $\text{FAR} < 6/\text{yr} = 1/2 \text{ months}$. For each candidate, we report the SNR, FAR and p_{astro} .	165
7.3 Sensitive spacetime volume estimated for the BGP population at a nominal FAR of 1/year and the FAR corresponding to the loudest event not reported in the GWTC. In the first case, we report the $\langle VT \rangle$ separately for O3a and O3b.	168

Acknowledgments

Here I am at the airport of Milan, waiting for a flight to Amsterdam, pretty much like four years ago when I took the first of many many planes from Milan to Amsterdam and viceversa. In these four years, so many things happened and so many new people joined my life that is hard to keep track of everything and to say thank you to everybody. I will try as hard as I can and apologies if I forget someone—it wasn't intentional.

First of all, thanks to Walter for teaching me a lot about gravitational waves and to spark my interest in research, which eventually brought me here—talking to you have always been a source of motivating. Of course, I'm so grateful to Sarah for providing me the great opportunity of moving to the Netherlands and starting this experience and also thanks for helping me to make the most out of it: whenever we managed to talk, you were always able to give me the confidence I was on the right track—that was precious. Also thanks to Chris are in order: for helping me to navigate in the Dutch academia and to help with me with new motivation in a period of struggle.

Thanks also to my Utrecht colleagues, Harsh, Thibeau, Tomek, Luca, Olaf, Rik, Bas, Quirijn and Justin, for sharing lunches, casual chats and lovely coffee breaks—it was fun to come to the office. Also thanks to Andrew and Bhooshan for taking the time to teach me a lot of new things. A special mention goes to Melissa and Tom, for some amazing dinners and a US trip which wouldn't have been so lovely without you.

Ever since the BND school in 2021, Nikhef have been a fun place to work from before going out in Amsterdam on Friday nights. On this regard, a special mention goes to Roberto and Mariia, for letting me invade their office. Besides a ton of fun Nikhef people, a special hug goes to Kelly, Pranati, Alex, (italian) Andrea, (spanish) Andrea, Marion, Giacomo, Tjip, Enrico, Carolina, Zeb, Ali, Guido, Petja and Evri. A good thanks to Pooya for our coffees in Utrecht, to Lluc for introducing me to INJI, to Thijs for an amazing graduation party which I will have to learn from and to Davide for making amazing pizzoccheri and for chatting on the trains to Maastricht. And of

course, very importantly, thanks to Carlo and Riccardo for forcing me into nights out in Amsterdam (without ever coming to Utrecht :D).

Going back to Utrecht, I couldn't be more grateful to the Sunday dinners with the "Italian PhDs", Anna, Martino and Sofia, which provided a safe space to complain about our job, the weather, the Netherlands and everything else. Thanks to the three of you: our dinners felt like family.

Ever since the Halloween party on my first night in the Netherlands, I met an amazing group of people, which grew over time and made me feel at home. Without the "Villagers", my life in the Netherlands wouldn't have been nearly as happy as it was. Thanks George for our dinners. Thanks also to Zog and Dani for staying in Utrecht: I hope not to disappoint you (too much) if I ever move to Amsterdam. Thanks Kostas for organizing a lovely trip to Berlin full of water and thanks to the amazing duo Melina and Athina for hosting parties and movies. Thanks Alvarito for the dinners at panuozzo and thanks Q for selecting the best melodic techno parties: I am happy I don't have to worry about this. Thanks Florian and Vera for coming to Viareggio and thanks to Chris for bringing me to my first bouldering experience after years without climbing. Thanks Ares, Alexander, Thomas and Rebecca for happening to be almost at every festival where I go. Thanks Abdelaki for bringing us to Egypt and thanks Lisa for biking together. Thanks Fabrizio for joining the Amsterdam dinners and introducing two Alessandro to me. Thanks Maite for ringing at my door asking for wine on that Halloween night where everything started. Thank you guys for making the last four years and endless series of fun.

Finally, thanks to my friends from Italy (but not only) who came to visit me in Utrecht: Marta, Marco, Pietro, Jack, Estelle, Lorenzo, Niya and Camilla. I will not name my other friends from Milan, who haven't visited me: please come, if you want to be mentioned on my next thesis. Thanks to mamma e papà for listening to everything and advising in the best way possible. Thanks to Chiara for reminding to me multiple times that what I do is not (completely) useless. Also thanks to Loredana, who was the first person to encourage me to move to the Netherlands. Of course, Marta, thanks for supporting me in everything and coming to Netherlands as much as you could: I know these four years of distance have been hard but I'm happy that you stayed with me regardless. Looking forward to start a life here together.

The people I have had around me have been the best part of my life in the Netherlands as well as the best part of my PhD. I couldn't be more grateful and I'm so excited to keep enjoying things together in the next years of my life in the Netherlands.

Bibliography

- [1] Stefano Schmidt, Bhooshan Gadre, and Sarah Caudill. Gravitational-wave template banks for novel compact binaries. *Phys. Rev. D*, 109(4):042005, 2024.
- [2] Stefano Schmidt and Sarah Caudill. Novel signal-consistency test for gravitational-wave searches of generic black hole binaries. *Phys. Rev. D*, 110(2):023042, 2024.
- [3] Stefano Schmidt et al. Searching for gravitational-wave signals from precessing black hole binaries with the GstLAL pipeline. *Phys. Rev. D*, 110(2):023038, 2024.
- [4] Stefano Schmidt et al. Searching for Asymmetric and Heavily Precessing Binary Black Holes in the Gravitational Wave Data from the LIGO Third Observing Run. *Phys. Rev. Lett.*, 133(20):201401, 2024.
- [5] Stefano Schmidt, Matteo Breschi, Rossella Gamba, Giulia Pagano, Piero Retegno, Gunnar Riemenschneider, Sebastiano Bernuzzi, Alessandro Nagar, and Walter Del Pozzo. Machine Learning Gravitational Waves from Binary Black Hole Mergers. *Phys. Rev. D*, 103(4):043020, 2021.
- [6] Jacopo Tissino, Gregorio Carullo, Matteo Breschi, Rossella Gamba, Stefano Schmidt, and Sebastiano Bernuzzi. Combining effective-one-body accuracy and reduced-order-quadrature speed for binary neutron star merger parameter estimation with machine learning. *Phys. Rev. D*, 107(8):084037, 2023.
- [7] Tim Grimbergen, Stefano Schmidt, Chinmay Kalaghatgi, and Chris van den Broeck. Generating higher order modes from binary black hole mergers with machine learning. *Phys. Rev. D*, 109(10):104065, 2024.
- [8] Paloma Laguarta et al. Detection of anomalies amongst LIGO’s glitch populations with autoencoders. *Class. Quant. Grav.*, 41(5):055004, 2024.

- [9] Alessandro Martini, Stefano Schmidt, Gregory Ashton, and Walter Del Pozzo. Maximum entropy spectral analysis: an application to gravitational waves data analysis. *Eur. Phys. J. C*, 84(10):1023, 2024.
- [10] Remo Ruffini and John A. Wheeler. Introducing the black hole. *Phys. Today*, 24(1):30, 1971.
- [11] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015.
- [12] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015.
- [13] T. Akutsu et al. Overview of KAGRA: Detector design and construction history. *PTEP*, 2021(5):05A101, 2021.
- [14] B. S. Sathyaprakash and B. F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Rev. Rel.*, 12:2, 2009.
- [15] B. P. Abbott et al. GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs. *Phys. Rev. X*, 9(3):031040, 2019.
- [16] R. Abbott et al. GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run. *Phys. Rev. X*, 11:021053, 2021.
- [17] R. Abbott et al. GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run. *Phys. Rev. D*, 109(2):022001, 2024.
- [18] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run. *Phys. Rev. X*, 13(4):041039, 2023.
- [19] B. P. Abbott et al. Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo. *Astrophys. J. Lett.*, 882(2):L24, 2019.
- [20] R. Abbott et al. Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog. *Astrophys. J. Lett.*, 913(1):L7, 2021.
- [21] R. Abbott et al. Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3. *Phys. Rev. X*, 13(1):011048, 2023.

- [22] B. P. Abbott et al. Estimating the Contribution of Dynamical Ejecta in the Kilonova Associated with GW170817. *Astrophys. J. Lett.*, 850(2):L39, 2017.
- [23] B. P. Abbott et al. Properties of the binary neutron star merger GW170817. *Phys. Rev. X*, 9(1):011001, 2019.
- [24] B. P. Abbott et al. GW170817: Measurements of neutron star radii and equation of state. *Phys. Rev. Lett.*, 121(16):161101, 2018.
- [25] Peter T. H. Pang et al. An updated nuclear-physics and multi-messenger astrophysics framework for binary neutron star mergers. *Nature Commun.*, 14(1):8352, 2023.
- [26] B. P. Abbott et al. A gravitational-wave standard siren measurement of the Hubble constant. *Nature*, 551(7678):85–88, 2017.
- [27] Salvatore Vitale and Hsin-Yu Chen. Measuring the Hubble constant with neutron star black hole mergers. *Phys. Rev. Lett.*, 121(2):021303, 2018.
- [28] B. P. Abbott et al. A Gravitational-wave Measurement of the Hubble Constant Following the Second Observing Run of Advanced LIGO and Virgo. *Astrophys. J.*, 909(2):218, 2021.
- [29] R. Abbott et al. Constraints on the Cosmic Expansion History from GWTC-3. *Astrophys. J.*, 949(2):76, 2023.
- [30] B. P. Abbott et al. Tests of general relativity with GW150914. *Phys. Rev. Lett.*, 116(22):221101, 2016. [Erratum: *Phys. Rev. Lett.* 121, 129902 (2018)].
- [31] B. P. Abbott et al. Tests of General Relativity with GW170817. *Phys. Rev. Lett.*, 123(1):011102, 2019.
- [32] 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):104036, 2019.
- [33] 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):122002, 2021.
- [34] R. Abbott et al. Tests of General Relativity with GWTC-3. 12 2021.
- [35] M. Punturo et al. The Einstein Telescope: A third-generation gravitational wave observatory. *Class. Quant. Grav.*, 27:194002, 2010.

- [36] Michele Maggiore et al. Science Case for the Einstein Telescope. *JCAP*, 03:050, 2020.
- [37] Matthew Evans et al. A Horizon Study for Cosmic Explorer: Science, Observatories, and Community. 9 2021.
- [38] Pau Amaro-Seoane et al. Laser Interferometer Space Antenna. 2 2017.
- [39] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin induced orbital precession and its modulation of the gravitational wave forms from merging binaries. *Phys. Rev. D*, 49:6274–6297, 1994.
- [40] Lawrence E. Kidder, Clifford M. Will, and Alan G. Wiseman. Spin effects in the inspiral of coalescing compact binaries. *Phys. Rev. D*, 47(10):R4183–R4187, 1993.
- [41] Lawrence E. Kidder. Coalescing binary systems of compact objects to post-Newtonian 5/2 order. 5. Spin effects. *Phys. Rev. D*, 52:821–847, 1995.
- [42] Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects: Adiabatic limit. *Phys. Rev. D*, 67:104025, 2003. [Erratum: *Phys. Rev. D* 74, 029904 (2006)].
- [43] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, Badri Krishnan, and David Merritt. Spin Flips and Precession in Black-Hole-Binary Mergers. *Phys. Rev. D*, 75:064030, 2007.
- [44] Etienne Racine, Alessandra Buonanno, and Lawrence E. Kidder. Recoil velocity at 2PN order for spinning black hole binaries. *Phys. Rev. D*, 80:044010, 2009.
- [45] Alejandro Bohe, Sylvain Marsat, Guillaume Faye, and Luc Blanchet. Next-to-next-to-leading order spin-orbit effects in the near-zone metric and precession equations of compact binaries. *Class. Quant. Grav.*, 30:075017, 2013.
- [46] Alejandro Bohé, Guillaume Faye, Sylvain Marsat, and Edward K. Porter. Quadratic-in-spin effects in the orbital dynamics and gravitational-wave energy flux of compact binaries at the 3PN order. *Class. Quant. Grav.*, 32(19):195010, 2015.
- [47] Geraint Pratten et al. Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes. *Phys. Rev. D*, 103(10):104056, 2021.

- [48] Héctor Estellés, Antoni Ramos-Buades, Sascha Husa, Cecilio García-Quirós, Marta Colleoni, Leïla Haegel, and Rafel Jaume. Phenomenological time domain model for dominant quadrupole gravitational wave signal of coalescing binary black holes. *Phys. Rev. D*, 103(12):124060, 2021.
- [49] Sarp Akcay, Rossella Gamba, and Sebastiano Bernuzzi. Hybrid post-Newtonian effective-one-body scheme for spin-precessing compact-binary waveforms up to merger. *Phys. Rev. D*, 103(2):024014, 2021.
- [50] Eleanor Hamilton, Lionel London, Jonathan E. Thompson, Edward Fauchon-Jones, Mark Hannam, Chinmay Kalaghatgi, Sebastian Khan, Francesco Pannarale, and Alex Vano-Vinuales. Model of gravitational waves from precessing black-hole binaries through merger and ringdown. *Phys. Rev. D*, 104(12):124027, 2021.
- [51] Rossella Gamba, Sarp Akçay, Sebastiano Bernuzzi, and Jake Williams. Effective-one-body waveforms for precessing coalescing compact binaries with post-Newtonian twist. *Phys. Rev. D*, 106(2):024020, 2022.
- [52] Antoni Ramos-Buades, Alessandra Buonanno, Héctor Estellés, Mohammed Khalil, Deyan P. Mihaylov, Serguei Ossokine, Lorenzo Pompili, and Mahlet Shiferaw. SEOBNRv5PHM: Next generation of accurate and efficient multipolar precessing-spin effective-one-body waveforms for binary black holes. 3 2023.
- [53] Charlie Hoy, Cameron Mills, and Stephen Fairhurst. Evidence for subdominant multipole moments and precession in merging black-hole-binaries from GWTC-2.1. *Phys. Rev. D*, 106(2):023019, 2022.
- [54] R. Abbott et al. GW190521: A Binary Black Hole Merger with a Total Mass of $150M_{\odot}$. *Phys. Rev. Lett.*, 125(10):101102, 2020.
- [55] Rachel C. Zhang, Giacomo Fragione, Chase Kimball, and Vicky Kalogera. On the Likely Dynamical Origin of GW191109 and Binary Black Hole Mergers with Negative Effective Spin. *Astrophys. J.*, 954(1):23, 2023.
- [56] Mark Hannam et al. General-relativistic precession in a black-hole binary. *Nature*, 610(7933):652–655, 2022.
- [57] Qianyun Yun, Wen-Biao Han, Qian Hu, and Haiguang Xu. Precessing binary black holes as better dark sirens. *Mon. Not. Roy. Astron. Soc.*, 527(1):L60–L65, 2023.

- [58] 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):L9, 2018.
- [59] Will M. Farr, Simon Stevenson, M. Coleman Miller, Ilya Mandel, Ben Farr, and Alberto Vecchio. Distinguishing Spin-Aligned and Isotropic Black Hole Populations With Gravitational Waves. *Nature*, 548:426, 2017.
- [60] 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. *Phys. Rev. D*, 106(2):023001, 2022.
- [61] Michela Mapelli. *Formation Channels of Single and Binary Stellar-Mass Black Holes*. 2021.
- [62] Salvatore Vitale, Sylvia Biscoveanu, and Colm Talbot. Spin it as you like: The (lack of a) measurement of the spin tilt distribution with LIGO-Virgo-KAGRA binary black holes. *Astron. Astrophys.*, 668:L2, 2022.
- [63] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D. E. Creighton. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys. Rev. D*, 85:122006, 2012.
- [64] Larne Pekowsky, James Healy, Deirdre Shoemaker, and Pablo Laguna. Impact of higher-order modes on the detection of binary black hole coalescences. *Phys. Rev. D*, 87(8):084008, 2013.
- [65] Stephen M. Privitera. *The Importance of Spin for Observing Gravitational Waves from Coalescing Compact Binaries with LIGO and Virgo*. Phd thesis, California Institute of Technology, Pasadena, CA, May 2014. Available at <https://resolver.caltech.edu/CaltechTHESIS:05282014-160218103>.
- [66] Tito Dal Canton, Andrew P. Lundgren, and Alex B. Nielsen. Impact of precession on aligned-spin searches for neutron-star–black-hole binaries. *Phys. Rev. D*, 91(6):062010, 2015.
- [67] Vijay Varma, Parameswaran Ajith, Sascha Husa, Juan Calderon Bustillo, Mark Hannam, and Michael Pürrer. Gravitational-wave observations of binary black holes: Effect of nonquadrupole modes. *Phys. Rev. D*, 90(12):124004, 2014.
- [68] Antoni Ramos-Buades, Shubhanshu Tiwari, Maria Haney, and Sascha Husa. Impact of eccentricity on the gravitational wave searches for binary black holes: High mass case. *Phys. Rev. D*, 102(4):043005, 2020.

- [69] Alexander H. Nitz, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Shichao Wu, Marlin Schäfer, Rahul Dhurkunde, and Collin D. Capano. 4-OGC: Catalog of Gravitational Waves from Compact Binary Mergers. *Astrophys. J.*, 946(2):59, 2023.
- [70] Alexander H. Nitz, Collin D. Capano, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Marlin Schäfer, Rahul Dhurkunde, and Miriam Cabero. 3-OGC: Catalog of Gravitational Waves from Compact-binary Mergers. *Astrophys. J.*, 922(1):76, 2021.
- [71] Seth Olsen, Tejaswi Venumadhav, Jonathan Mushkin, Javier Roulet, Barak Zackay, and Matias Zaldarriaga. New binary black hole mergers in the LIGO-Virgo O3a data. *Phys. Rev. D*, 106(4):043009, 2022.
- [72] Ajit Kumar Mehta, Seth Olsen, Digvijay Wadekar, Javier Roulet, Tejaswi Venumadhav, Jonathan Mushkin, Barak Zackay, and Matias Zaldarriaga. New binary black hole mergers in the LIGO-Virgo O3b data. 11 2023.
- [73] 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):104038, 2017.
- [74] Stephen Fairhurst, Rhys Green, Mark Hannam, and Charlie Hoy. When will we observe binary black holes precessing? *Phys. Rev. D*, 102(4):041302, 2020.
- [75] Rahul Dhurkunde and Alexander H. Nitz. Sensitivity of spin-aligned searches for neutron star-black hole systems using future detectors. *Phys. Rev. D*, 106:103035, Nov 2022.
- [76] Charlie Hoy, Stephen Fairhurst, and Ilya Mandel. Precession and higher order multipoles in binary black holes (and lack thereof). 8 2024.
- [77] Cody Messick et al. Analysis Framework for the Prompt Discovery of Compact Binary Mergers in Gravitational-wave Data. *Phys. Rev. D*, 95(4):042001, 2017.
- [78] Surabhi Sachdev et al. The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs. 1 2019.
- [79] Chad Hanna et al. Fast evaluation of multidecoder consistency for real-time gravitational wave searches. *Phys. Rev. D*, 101(2):022003, 2020.
- [80] Kipp Cannon, Sarah Caudill, Chiwai Chan, Bryce Cousins, Jolien D.E. Creighton, Becca Ewing, Heather Fong, Patrick Godwin, Chad Hanna, Shaun

Hooper, Rachael Huxford, Ryan Magee, Duncan Meacher, Cody Messick, Soichiro Morisaki, Debnandini Mukherjee, Hiroaki Ohta, Alexander Pace, Stephen Privitera, Iris de Ruiter, Surabhi Sachdev, Leo Singer, Divya Singh, Ron Tapia, Leo Tsukada, Daichi Tsuna, Takuya Tsutsui, Koh Ueno, Aaron Viets, Leslie Wade, and Madeline Wade. Gstlal: A software framework for gravitational wave discovery. *SoftwareX*, 14:100680, 2021.

[81] Becca Ewing et al. Performance of the low-latency GstLAL inspiral search towards LIGO, Virgo, and KAGRA's fourth observing run. *Phys. Rev. D*, 109(4):042008, 2024.

[82] Leo Tsukada et al. Improved ranking statistics of the GstLAL inspiral search for compact binary coalescences. *Phys. Rev. D*, 108(4):043004, 2023.

[83] Rich Abbott et al. Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo. *SoftwareX*, 13:100658, 2021.

[84] R. Abbott et al. Open Data from the Third Observing Run of LIGO, Virgo, KAGRA, and GEO. *Astrophys. J. Suppl.*, 267(2):29, 2023.

[85] Michele Maggiore. *Gravitational Waves. Vol. 1: Theory and Experiments*. Oxford Master Series in Physics. Oxford University Press, 2007.

[86] Albert Einstein. The foundation of the general theory of relativity. *Annalen Phys.*, 49(7):769–822, 1916.

[87] Harald Dimmelmeier, Christian D. Ott, Andreas Marek, and H. Thomas Janka. The Gravitational Wave Burst Signal from Core Collapse of Rotating Stars. *Phys. Rev. D*, 78:064056, 2008.

[88] Ernazar Abdikamalov, Giulia Pagliaroli, and David Radice. Gravitational Waves from Core-Collapse Supernovae. 10 2020.

[89] David Vartanyan, Adam Burrows, Tianshu Wang, Matthew S. B. Coleman, and Christopher J. White. Gravitational-wave signature of core-collapse supernovae. *Phys. Rev. D*, 107(10):103015, 2023.

[90] T. W. B. Kibble. Topology of Cosmic Domains and Strings. *J. Phys. A*, 9:1387–1398, 1976.

[91] Tanmay Vachaspati and Alexander Vilenkin. Gravitational Radiation from Cosmic Strings. *Phys. Rev. D*, 31:3052, 1985.

[92] M. Sakellariadou. Gravitational waves emitted from infinite strings. *Phys. Rev. D*, 42:354–360, 1990. [Erratum: Phys.Rev.D 43, 4150 (1991)].

- [93] Chia-Feng Chang and Yanou Cui. Gravitational waves from global cosmic strings and cosmic archaeology. *JHEP*, 03:114, 2022.
- [94] Reinhard Prix. *Gravitational Waves from Spinning Neutron Stars*, pages 651–685. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [95] Paul D. Lasky. Gravitational Waves from Neutron Stars: A Review. *Publ. Astron. Soc. Austral.*, 32:e034, 2015.
- [96] Kostas Glampedakis and Leonardo Gualtieri. *Gravitational Waves from Single Neutron Stars: An Advanced Detector Era Survey*, pages 673–736. Springer International Publishing, Cham, 2018.
- [97] Fabian Gittins. Gravitational waves from neutron-star mountains. *Class. Quant. Grav.*, 41(4):043001, 2024.
- [98] Nelson Christensen. Stochastic Gravitational Wave Backgrounds. *Rept. Prog. Phys.*, 82(1):016903, 2019.
- [99] Tania Regimbau. The astrophysical gravitational wave stochastic background. *Res. Astron. Astrophys.*, 11:369–390, 2011.
- [100] Chiara Caprini and Daniel G. Figueira. Cosmological Backgrounds of Gravitational Waves. *Class. Quant. Grav.*, 35(16):163001, 2018.
- [101] Sean M. Carroll. Lecture notes on general relativity. 12 1997.
- [102] K. S. Thorne. Multipole Expansions of Gravitational Radiation. *Rev. Mod. Phys.*, 52:299–339, 1980.
- [103] D. Brown, S. Fairhurst, B. Krishnan, R. A. Mercer, R. K. Koppaparu, L. Santa-maria, and J. T. Whelan. Data formats for numerical relativity waves. 9 2007.
- [104] Lawrence E. Kidder. Using full information when computing modes of post-Newtonian waveforms from inspiralling compact binaries in circular orbit. *Phys. Rev. D*, 77:044016, 2008.
- [105] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Siddhartha Sinha. The Third post-Newtonian gravitational wave polarisations and associated spherical harmonic modes for inspiralling compact binaries in quasi-circular orbits. *Class. Quant. Grav.*, 25:165003, 2008. [Erratum: *Class.Quant.Grav.* 29, 239501 (2012)].
- [106] R. Abbott et al. GW190412: Observation of a Binary-Black-Hole Coalescence with Asymmetric Masses. *Phys. Rev. D*, 102(4):043015, 2020.

- [107] R. Abbott et al. GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object. *Astrophys. J. Lett.*, 896(2):L44, 2020.
- [108] Karl Schwarzschild. On the gravitational field of a mass point according to Einstein's theory. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1916:189–196, 1916.
- [109] Roy P. Kerr. Gravitational field of a spinning mass as an example of algebraically special metrics. *Phys. Rev. Lett.*, 11:237–238, 1963.
- [110] Saul A. Teukolsky. The Kerr Metric. *Class. Quant. Grav.*, 32(12):124006, 2015.
- [111] H. Reissner. Über die Eigengravitation des elektrischen Feldes nach der Einsteinschen Theorie. *Annalen Phys.*, 355(9):106–120, 1916.
- [112] E T. Newman, R. Couch, K. Chinnapared, A. Exton, A. Prakash, and R. Torrence. Metric of a Rotating, Charged Mass. *J. Math. Phys.*, 6:918–919, 1965.
- [113] Tim Adamo and E. T. Newman. The Kerr-Newman metric: A Review. *Scholarpedia*, 9:31791, 2014.
- [114] Robert M. Wald. *General Relativity*. Chicago Univ. Pr., Chicago, USA, 1984.
- [115] S. L. Shapiro and S. A. Teukolsky. *Black holes, white dwarfs, and neutron stars: The physics of compact objects*. 1983.
- [116] Subrahmanyan Chandrasekhar. The maximum mass of ideal white dwarfs. *Astrophys. J.*, 74:81–82, 1931.
- [117] Vassiliki Kalogera and Gordon Baym. The maximum mass of a neutron star. *Astrophys. J. Lett.*, 470:L61–L64, 1996.
- [118] Daniel A. Godzieba, David Radice, and Sebastiano Bernuzzi. On the maximum mass of neutron stars and GW190814. *Astrophys. J.*, 908(2):122, 2021.
- [119] Félix Mirabel. The Formation of Stellar Black Holes. *New Astron. Rev.*, 78:1–15, 2017.
- [120] Laura Ferrarese and David Merritt. Supermassive black holes. *Phys. World*, 15N6:41–46, 2002.
- [121] Ramesh Narayan and Jeffrey E. McClintock. Observational Evidence for Black Holes. 12 2013.

- [122] Laura Ferrarese and David Merritt. A Fundamental relation between supermassive black holes and their host galaxies. *Astrophys. J. Lett.*, 539:L9, 2000.
- [123] Laura Ferrarese and Holland Ford. Supermassive black holes in galactic nuclei: Past, present and future research. *Space Sci. Rev.*, 116:523–624, 2005.
- [124] Aaron Smith and Volker Bromm. Supermassive Black Holes in the Early Universe. *Contemp. Phys.*, 60(2):111–126, 2019.
- [125] Marta Volonteri, Melanie Habouzit, and Monica Colpi. The origins of massive black holes. *Nature Rev. Phys.*, 3(11):732–743, 2021.
- [126] Krzysztof Belczynski, Daniel E. Holz, Tomasz Bulik, and Richard O’Shaughnessy. The first gravitational-wave source from the isolated evolution of two 40-100 Msun stars. *Nature*, 534:512, 2016.
- [127] Hans A. Bethe and G. E. Brown. Evolution of binary compact objects which merge. *Astrophys. J.*, 506:780–789, 1998.
- [128] Michal Dominik, Krzysztof Belczynski, Christopher Fryer, Daniel Holz, Emanuele Berti, Tomasz Bulik, Ilya Mandel, and Richard O’Shaughnessy. Double Compact Objects I: The Significance of the Common Envelope on Merger Rates. *Astrophys. J.*, 759:52, 2012.
- [129] Michal Dominik, Emanuele Berti, Richard O’Shaughnessy, Ilya Mandel, Krzysztof Belczynski, Christopher Fryer, Daniel E. Holz, Tomasz Bulik, and Francesco Pannarale. Double Compact Objects III: Gravitational Wave Detection Rates. *Astrophys. J.*, 806(2):263, 2015.
- [130] Ilya Mandel and Selma E. de Mink. Merging binary black holes formed through chemically homogeneous evolution in short-period stellar binaries. *Mon. Not. Roy. Astron. Soc.*, 458(3):2634–2647, 2016.
- [131] Kohei Inayoshi, Ryosuke Hirai, Tomoya Kinugawa, and Kenta Hotokezaka. Formation pathway of Population III coalescing binary black holes through stable mass transfer. *Mon. Not. Roy. Astron. Soc.*, 468(4):5020–5032, 2017.
- [132] Shrinivas F. Kulkarni, Steve McMillan, and Piet Hut. Stellar black holes in globular clusters. *Nature*, 364:421, 1993.
- [133] Brunetto Marco Ziosi, Michela Mapelli, Marica Branchesi, and Giuseppe Torsten. Dynamics of stellar black holes in young star clusters with different metallicities – II. Black hole-black hole binaries. *Mon. Not. Roy. Astron. Soc.*, 441(4):3703–3717, 2014.

- [134] Carl L. Rodriguez, Meagan Morscher, Bharath Pattaibiraman, Sourav Chatterjee, Carl-Johan Haster, and Frederic A. Rasio. Binary Black Hole Mergers from Globular Clusters: Implications for Advanced LIGO. *Phys. Rev. Lett.*, 115(5):051101, 2015. [Erratum: Phys.Rev.Lett. 116, 029901 (2016)].
- [135] Fabio Antonini, Silvia Toonen, and Adrian S. Hamers. Binary black hole mergers from field triples: properties, rates and the impact of stellar evolution. *Astrophys. J.*, 841(2):77, 2017.
- [136] Giacomo Fragione, Evgeni Grishin, Nathan W. C. Leigh, Hagai. B. Perets, and Rosalba Perna. Black hole and neutron star mergers in galactic nuclei. *Mon. Not. Roy. Astron. Soc.*, 488(1):47–63, 2019.
- [137] Giacomo Fragione, Nathan Leigh, and Rosalba Perna. Black hole and neutron star mergers in Galactic Nuclei: the role of triples. *Mon. Not. Roy. Astron. Soc.*, 488(2):2825–2835, 2019.
- [138] Imre Bartos, Bence Kocsis, Zolt Haiman, and Szabolcs Márka. Rapid and Bright Stellar-mass Binary Black Hole Mergers in Active Galactic Nuclei. *Astrophys. J.*, 835(2):165, 2017.
- [139] Carl L. Rodriguez, Michael Zevin, Chris Pankow, Vasiliki Kalogera, and Frederic A. Rasio. Illuminating Black Hole Binary Formation Channels with Spins in Advanced LIGO. *Astrophys. J. Lett.*, 832(1):L2, 2016.
- [140] P. C. Peters and J. Mathews. Gravitational radiation from point masses in a Keplerian orbit. *Phys. Rev.*, 131:435–439, 1963.
- [141] Jolien Creighton and Warren Anderson. *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis*. 2011.
- [142] Luc Blanchet. Gravitational Radiation from Post-Newtonian Sources and In-spiralling Compact Binaries. *Living Rev. Rel.*, 17:2, 2014.
- [143] Alessandra Buonanno, Bala Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-Newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev. D*, 80:084043, 2009.
- [144] Frans Pretorius. Numerical relativity using a generalized harmonic decomposition. *Class. Quant. Grav.*, 22:425–452, 2005.
- [145] Frans Pretorius. Binary Black Hole Coalescence. 8 2007.
- [146] Marcus Ansorg, Bernd Bruegmann, and Wolfgang Tichy. A Single-domain spectral method for black hole puncture data. *Phys. Rev. D*, 70:064011, 2004.

- [147] Luca Baiotti, Ian Hawke, Pedro J. Montero, Frank Löffler, Luciano Rezzolla, Nikolaos Stergioulas, Jose A. Font, and Ed Seidel. Three-dimensional relativistic simulations of rotating neutron star collapse to a Kerr black hole. *Phys. Rev. D*, 71:024035, 2005.
- [148] Tim Dietrich, Niclas Moldenhauer, Nathan K. Johnson-McDaniel, Sebastiano Bernuzzi, Charalampos M. Markakis, Bernd Brügmann, and Wolfgang Tichy. Binary Neutron Stars with Generic Spin, Eccentricity, Mass ratio, and Compactness - Quasi-equilibrium Sequences and First Evolutions. *Phys. Rev. D*, 92(12):124007, 2015.
- [149] Kyriakos Destounis and Francisco Duque. Black-hole spectroscopy: quasinormal modes, ringdown stability and the pseudospectrum. 8 2023.
- [150] Alessandra Buonanno and Thibault Damour. Transition from inspiral to plunge in binary black hole coalescences. *Phys. Rev.*, D62:064015, 2000.
- [151] Thibault Damour and Alessandro Nagar. An improved analytical description of inspiralling and coalescing black-hole binaries. *Phys. Rev.*, D79:081503, 2009.
- [152] Roberto Cotesta, Alessandra Buonanno, Alejandro Bohé, Andrea Taracchini, Ian Hinder, and Serguei Ossokine. Enriching the Symphony of Gravitational Waves from Binary Black Holes by Tuning Higher Harmonics. *Phys. Rev. D*, 98(8):084028, 2018.
- [153] Alessandro Nagar, Gunnar Riemenschneider, Geraint Pratten, Piero Rettegno, and Francesco Messina. A multipolar effective one body waveform model for spin-aligned black hole binaries. 2020.
- [154] Danilo Chiaramello and Alessandro Nagar. Faithful analytical effective-one-body waveform model for spin-aligned, moderately eccentric, coalescing black hole binaries. *Phys. Rev. D*, 101(10):101501, 2020.
- [155] Serguei Ossokine et al. Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation. *Phys. Rev. D*, 102(4):044055, 2020.
- [156] Alessandro Nagar, Alice Bonino, and Piero Rettegno. Effective one-body multipolar waveform model for spin-aligned, quasicircular, eccentric, hyperbolic black hole binaries. *Phys. Rev. D*, 103(10):104021, 2021.
- [157] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational

waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era. *Phys. Rev. D*, 93(4):044007, 2016.

[158] Geraint Pratten, Sascha Husa, Cecilio García-Quiros, Marta Colleoni, Antoni Ramos-Buades, Hector Estelles, and Rafel Jaume. Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes. *Phys. Rev. D*, 102(6):064001, 2020.

[159] Héctor Estellés, Sascha Husa, Marta Colleoni, David Keitel, Maite Mateu-Lucena, Cecilio García-Quirós, Antoni Ramos-Buades, and Angela Borchers. Time-domain phenomenological model of gravitational-wave subdominant harmonics for quasicircular nonprecessing binary black hole coalescences. *Phys. Rev. D*, 105(8):084039, 2022.

[160] Mohammed Khalil, Alessandra Buonanno, Jan Steinhoff, and Justin Vines. Radiation-reaction force and multipolar waveforms for eccentric, spin-aligned binaries in the effective-one-body formalism. *Phys. Rev. D*, 104(2):024046, 2021.

[161] Antoni Ramos-Buades, Alessandra Buonanno, Mohammed Khalil, and Serguei Ossokine. Effective-one-body multipolar waveforms for eccentric binary black holes with nonprecessing spins. *Phys. Rev. D*, 105(4):044035, 2022.

[162] Rossella Gamba, Matteo Breschi, Gregorio Carullo, Simone Albanesi, Piero Rettegno, Sebastiano Bernuzzi, and Alessandro Nagar. GW190521 as a dynamical capture of two nonspinning black holes. *Nature Astron.*, 7(1):11–17, 2023.

[163] Teagan A. Clarke, Isobel M. Romero-Shaw, Paul D. Lasky, and Eric Thrane. Gravitational-wave inference for eccentric binaries: the argument of periapsis. *Mon. Not. Roy. Astron. Soc.*, 517(3):3778–3784, 2022.

[164] Alice Bonino, Rossella Gamba, Patricia Schmidt, Alessandro Nagar, Geraint Pratten, Matteo Breschi, Piero Rettegno, and Sebastiano Bernuzzi. Inferring eccentricity evolution from observations of coalescing binary black holes. *Phys. Rev. D*, 107(6):064024, 2023.

[165] Antoni Ramos-Buades, Alessandra Buonanno, and Jonathan Gair. Bayesian inference of binary black holes with inspiral-merger-ringdown waveforms using two eccentric parameters. *Phys. Rev. D*, 108(12):124063, 2023.

[166] Souradeep Pal and K. Rajesh Nayak. Swarm-intelligent search for gravitational waves from eccentric binary mergers. 7 2023.

- [167] Michael Kesden, Davide Gerosa, Richard O’Shaughnessy, Emanuele Berti, and Ulrich Sperhake. Effective potentials and morphological transitions for binary black-hole spin precession. *Phys. Rev. Lett.*, 114(8):081103, 2015.
- [168] Davide Gerosa, Michael Kesden, Ulrich Sperhake, Emanuele Berti, and Richard O’Shaughnessy. Multi-timescale analysis of phase transitions in precessing black-hole binaries. *Phys. Rev. D*, 92:064016, 2015.
- [169] Davide Gerosa and Michael Kesden. PRECESSION: Dynamics of spinning black-hole binaries with python. *Phys. Rev. D*, 93(12):124066, 2016.
- [170] Davide Gerosa, Giulia Fumagalli, Matthew Mould, Giovanni Cavallotto, Diego Padilla Monroy, Daria Gangardt, and Viola De Renzis. Efficient multi-timescale dynamics of precessing black-hole binaries. *Phys. Rev. D*, 108(2):024042, 2023.
- [171] Patricia Schmidt, Mark Hannam, Sascha Husa, and P. Ajith. Tracking the precession of compact binaries from their gravitational-wave signal. *Phys. Rev. D*, 84:024046, 2011.
- [172] Patricia Schmidt, Mark Hannam, and Sascha Husa. Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and non-precessing inspiral signals. *Phys. Rev. D*, 86:104063, 2012.
- [173] Eric Poisson. Gravitational waves from inspiraling compact binaries: The Quadrupole moment term. *Phys. Rev. D*, 57:5287–5290, 1998.
- [174] Laszlo A. Gergely. Spin spin effects in radiating compact binaries. *Phys. Rev. D*, 61:024035, 2000.
- [175] Luc Blanchet, Guillaume Faye, Bala R. Iyer, and Benoit Joguet. Gravitational wave inspiral of compact binary systems to 7/2 postNewtonian order. *Phys. Rev. D*, 65:061501, 2002. [Erratum: Phys.Rev.D 71, 129902 (2005)].
- [176] Luc Blanchet, Thibault Damour, Gilles Esposito-Farese, and Bala R. Iyer. Gravitational radiation from inspiralling compact binaries completed at the third post-Newtonian order. *Phys. Rev. Lett.*, 93:091101, 2004.
- [177] Luc Blanchet, Alessandra Buonanno, and Guillaume Faye. Higher-order spin effects in the dynamics of compact binaries. II. Radiation field. *Phys. Rev. D*, 74:104034, 2006. [Erratum: Phys.Rev.D 75, 049903 (2007), Erratum: Phys.Rev.D 81, 089901 (2010)].

- [178] Guillaume Faye, Luc Blanchet, and Alessandra Buonanno. Higher-order spin effects in the dynamics of compact binaries. I. Equations of motion. *Phys. Rev. D*, 74:104033, 2006.
- [179] Michael Boyle, Robert Owen, and Harald P. Pfeiffer. A geometric approach to the precession of compact binaries. *Phys. Rev. D*, 84:124011, 2011.
- [180] Mark Hannam, Patricia Schmidt, Alejandro Bohé, Leïla Haegel, Sascha Husa, Frank Ohme, Geraint Pratten, and Michael Pürer. Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms. *Phys. Rev. Lett.*, 113(15):151101, 2014.
- [181] Michael Boyle, Lawrence E. Kidder, Serguei Ossokine, and Harald P. Pfeiffer. Gravitational-wave modes from precessing black-hole binaries. 9 2014.
- [182] Rhys Green, Charlie Hoy, Stephen Fairhurst, Mark Hannam, Francesco Pannarale, and Cory Thomas. Identifying when Precession can be Measured in Gravitational Waveforms. *Phys. Rev. D*, 103(12):124023, 2021.
- [183] 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):024043, 2015.
- [184] Benjamin Farr, Evan Ochsner, Will M. Farr, and Richard O’Shaughnessy. A more effective coordinate system for parameter estimation of precessing compact binaries from gravitational waves. *Phys. Rev. D*, 90(2):024018, 2014.
- [185] Stephen Fairhurst, Rhys Green, Charlie Hoy, Mark Hannam, and Alistair Muir. Two-harmonic approximation for gravitational waveforms from precessing binaries. *Phys. Rev. D*, 102(2):024055, 2020.
- [186] Lucy M. Thomas, Patricia Schmidt, and Geraint Pratten. New effective precession spin for modeling multimodal gravitational waveforms in the strong-field regime. *Phys. Rev. D*, 103(8):083022, 2021.
- [187] Davide Gerosa, Matthew Mould, Daria Gangardt, Patricia Schmidt, Geraint Pratten, and Lucy M. Thomas. A generalized precession parameter χ_p to interpret gravitational-wave data. *Phys. Rev. D*, 103(6):064067, 2021.
- [188] B. P. Abbott et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.
- [189] Odylio Denys Aguiar. The Past, Present and Future of the Resonant-Mass Gravitational Wave Detectors. *Res. Astron. Astrophys.*, 11:1–42, 2011.

- [190] C. Fabry and A. Pérot. Théorie et applications d'une nouvelle méthode de spectroscopie interférentielle. *Annales de Chimie et de Physique*, 16(6):115–144, 1899.
- [191] B. P. Abbott et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. *Phys. Rev. D*, 95(6):062003, 2017.
- [192] Craig Cahillane, Joe Betzwieser, Duncan A. Brown, Evan Goetz, Evan D. Hall, Kiwamu Izumi, Shivaraj Kandhasamy, Sudarshan Karki, Jeff S. Kissel, Greg Mendell, Richard L. Savage, Darkhan Tuyenbayev, Alex Urban, Aaron Viets, Madeline Wade, and Alan J. Weinstein. Calibration uncertainty for advanced ligo's first and second observing runs. *Phys. Rev. D*, 96:102001, Nov 2017.
- [193] Ling Sun et al. Characterization of systematic error in Advanced LIGO calibration. *Class. Quant. Grav.*, 37(22):225008, 2020.
- [194] Madeline Wade, Aaron D. Viets, Theresa Chmiel, Madeline Stover, and Leslie Wade. Improving LIGO calibration accuracy by using time-dependent filters to compensate for temporal variations. *Class. Quant. Grav.*, 40(3):035001, 2023.
- [195] Andreas Freise and Kenneth Strain. Interferometer Techniques for Gravitational-Wave Detection. *Living Rev. Rel.*, 13:1, 2010.
- [196] Benjamin P. Abbott et al. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 93(11):112004, 2016. [Addendum: Phys.Rev.D 97, 059901 (2018)].
- [197] A. Buikema et al. Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Phys. Rev. D*, 102(6):062003, Sep 2020.
- [198] S. Braccini. The VIRGO suspensions. *Class. Quant. Grav.*, 19:1623–1629, 2002.
- [199] S. M. Aston et al. Update on quadruple suspension design for Advanced LIGO. *Class. Quant. Grav.*, 29(23):235004, 2012.
- [200] Jan Harms, Luca Naticchioni, Enrico Calloni, Rosario De Rosa, Fulvio Ricci, and Domenico D'Urso. A lower limit for Newtonian-noise models of the Einstein Telescope. *Eur. Phys. J. Plus*, 137(6):687, 2022.
- [201] Kamiel Janssens, Guillaume Boileau, Nelson Christensen, Francesca Badaracco, and Nick van Remortel. Impact of correlated seismic and correlated Newtonian noise on the Einstein Telescope. *Phys. Rev. D*, 106(4):042008, 2022.

Bibliography

- [202] Jan Harms and Ho Jung Paik. Newtonian-noise cancellation in full-tensor gravitational-wave detectors. *Phys. Rev. D*, 92(2):022001, 2015.
- [203] Soumen Koley et al. Design and implementation of a seismic Newtonian noise cancellation system for the Virgo gravitational-wave detector. *Eur. Phys. J. Plus*, 139(1):48, 2024.
- [204] Torrey T. Lyons, Martin W. Regehr, and Frederick J. Raab. Shot noise in gravitational-wave detectors with fabry–perot arms. *Appl. Opt.*, 39(36):6761–6770, Dec 2000.
- [205] J. Abadie et al. A Gravitational wave observatory operating beyond the quantum shot-noise limit: Squeezed light in application. *Nature Phys.*, 7:962–965, 2011.
- [206] J. Aasi et al. Enhancing the sensitivity of the LIGO gravitational wave detector by using squeezed states of light. *Nature Photon.*, 7:613–619, 2013.
- [207] L. McCuller et al. Frequency-Dependent Squeezing for Advanced LIGO. *Phys. Rev. Lett.*, 124(17):171102, 2020.
- [208] Nelson Christensen, Peter Shawhan, and Gabriela Gonzalez. Vetoos for inspiral triggers in LIGO data. *Class. Quant. Grav.*, 21:S1747–S1756, 2004.
- [209] B P Abbott et al. Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO’s first observing run. *Class. Quant. Grav.*, 35(6):065010, 2018.
- [210] R. Abbott et al. Data quality vetoes applied to the analysis of ligo data from the third observing run. <https://dcc.ligo.org/LIGO-T2100045/public>.
- [211] L Blackburn et al. The LSC Glitch Group: Monitoring Noise Transients during the fifth LIGO Science Run. *Class. Quant. Grav.*, 25:184004, 2008.
- [212] Michael Zevin et al. Gravity Spy: Integrating Advanced LIGO Detector Characterization, Machine Learning, and Citizen Science. *Class. Quant. Grav.*, 34(6):064003, 2017.
- [213] B. P. Abbott et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *Class. Quant. Grav.*, 33(13):134001, 2016.
- [214] Derek Davis et al. LIGO detector characterization in the second and third observing runs. *Class. Quant. Grav.*, 38(13):135014, 2021.

- [215] S. Soni et al. Reducing scattered light in LIGO’s third observing run. *Class. Quant. Grav.*, 38(2):025016, 2020.
- [216] Miriam Cabero et al. Blip glitches in Advanced LIGO data. *Class. Quant. Grav.*, 36(15):15, 2019.
- [217] B. P. Abbott et al. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.
- [218] Chris Pankow et al. Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817. *Phys. Rev. D*, 98(8):084016, 2018.
- [219] Daniel George, Hongyu Shen, and E. A. Huerta. Deep Transfer Learning: A new deep learning glitch classification method for advanced LIGO. 6 2017.
- [220] Jade Powell. Parameter Estimation and Model Selection of Gravitational Wave Signals Contaminated by Transient Detector Noise Glitches. *Class. Quant. Grav.*, 35(15):155017, 2018.
- [221] P. B. Covas et al. Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO. *Phys. Rev. D*, 97(8):082002, 2018.
- [222] J. C. Driggers et al. Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO. *Phys. Rev. D*, 99(4):042001, 2019.
- [223] D. Davis, T. J. Massinger, A. P. Lundgren, J. C. Driggers, A. L. Urban, and L. K. Nuttall. Improving the Sensitivity of Advanced LIGO Using Noise Subtraction. *Class. Quant. Grav.*, 36(5):055011, 2019.
- [224] Benjamin Steltner, Maria Alessandra Papa, and Heinz-Bernd Eggenstein. Identification and removal of non-Gaussian noise transients for gravitational-wave searches. *Phys. Rev. D*, 105(2):022005, 2022.
- [225] Katerina Chatzioannou, Neil Cornish, Marcella Wijngaarden, and Tyson B. Littenberg. Modeling compact binary signals and instrumental glitches in gravitational wave data. *Phys. Rev. D*, 103(4):044013, 2021.
- [226] J. Glanzer et al. Data quality up to the third observing run of advanced LIGO: Gravity Spy glitch classifications. *Class. Quant. Grav.*, 40(6):065004, 2023.
- [227] Melissa Lopez, Vincent Boudart, Kerwin Buijsman, Amit Reza, and Sarah Caudill. Simulating transient noise bursts in ligo with generative adversarial networks. *Phys. Rev. D*, 106:023027, Jul 2022.

- [228] Melissa Lopez, Vincent Boudart, Stefano Schmidt, and Sarah Caudill. Simulating transient noise bursts in ligo with gengli, 2022.
- [229] Tabata Aira Ferreira and Cesar Augusto Costa. Comparison between t-SNE and cosine similarity for LIGO glitches analysis. *Class. Quant. Grav.*, 39(16):165013, 2022.
- [230] Michael Zevin et al. Gravity Spy: lessons learned and a path forward. *Eur. Phys. J. Plus*, 139(1):100, 2024.
- [231] Soumya D. Mohanty and Mohammad A. T. Chowdhury. Glitch subtraction from gravitational wave data using adaptive spline fitting. *Class. Quant. Grav.*, 40(12):125001, 2023.
- [232] Chris Chatfield. *The analysis of time series: an introduction*. CRC Press, Florida, US, 6th edition, 2004.
- [233] P. Gregory. *Multivariate Gaussian from maximum entropy*, page 450–454. Cambridge University Press, 2005.
- [234] Gabriele Vajente, Yiwen Huang, Maximiliano Isi, Jenne C. Driggers, Jeffrey S. Kissel, Marek J. Szczerpanczyk, and Salvatore Vitale. Machine-learning nonstationary noise out of gravitational-wave detectors. *Phys. Rev. D*, 101(4):042003, 2020.
- [235] Barak Zackay, Tejaswi Venumadhav, Javier Roulet, Liang Dai, and Matias Zaldarriaga. Detecting gravitational waves in data with non-stationary and non-Gaussian noise. *Phys. Rev. D*, 104(6):063034, 2021.
- [236] Oliver Edy, Andrew Lundgren, and Laura K. Nuttall. Issues of mismodeling gravitational-wave data for parameter estimation. *Phys. Rev. D*, 103(12):124061, 2021.
- [237] Sumit Kumar, Alexander H. Nitz, and Xisco Jiménez Forteza. Parameter estimation with non stationary noise in gravitational waves data. 2 2022.
- [238] Oliver Edy. *Non-Stationarity in Gravitational-Wave Analysis*. Phd thesis, University of Portsmouth, Portsmouth, UK, September 2022. Available at <https://researchportal.port.ac.uk/en/studentTheses/non-stationarity-in-gravitational-wave-analysis>.
- [239] Peter Welch. The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, 15(2):70–73, 1967.

- [240] J.P. Burg. *Maximum Entropy Spectral Analysis*. Stanford Exploration Project. Stanford University, 1975.
- [241] R. Abbott et al. All-sky search for continuous gravitational waves from isolated neutron stars in the early O3 LIGO data. *Phys. Rev. D*, 104(8):082004, 2021.
- [242] R. Abbott et al. Search for continuous gravitational waves from 20 accreting millisecond x-ray pulsars in O3 LIGO data. *Phys. Rev. D*, 105:022002, 2022.
- [243] R. Abbott et al. Narrowband Searches for Continuous and Long-duration Transient Gravitational Waves from Known Pulsars in the LIGO-Virgo Third Observing Run. *Astrophys. J.*, 932(2):133, 2022.
- [244] R. Abbott et al. All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data. *Phys. Rev. D*, 106(10):102008, 2022.
- [245] R. Abbott et al. Search for continuous gravitational wave emission from the Milky Way center in O3 LIGO-Virgo data. *Phys. Rev. D*, 106(4):042003, 2022.
- [246] B. P. Abbott et al. A First Targeted Search for Gravitational-Wave Bursts from Core-Collapse Supernovae in Data of First-Generation Laser Interferometer Detectors. *Phys. Rev. D*, 94(10):102001, 2016.
- [247] B. P. Abbott et al. Optically targeted search for gravitational waves emitted by core-collapse supernovae during the first and second observing runs of advanced LIGO and advanced Virgo. *Phys. Rev. D*, 101(8):084002, 2020.
- [248] B. P. Abbott et al. Searches for Continuous Gravitational Waves from 15 Supernova Remnants and Fomalhaut b with Advanced LIGO. *Astrophys. J.*, 875(2):122, 2019. [Erratum: *Astrophys.J.* 918, 91 (2021)].
- [249] R. Abbott et al. Searches for Continuous Gravitational Waves from Young Supernova Remnants in the Early Third Observing Run of Advanced LIGO and Virgo. *Astrophys. J.*, 921(1):80, 2021.
- [250] R. Abbott et al. Search of the early O3 LIGO data for continuous gravitational waves from the Cassiopeia A and Vela Jr. supernova remnants. *Phys. Rev. D*, 105(8):082005, 2022.
- [251] Benjamin P. Abbott et al. Upper Limits on the Stochastic Gravitational-Wave Background from Advanced LIGO's First Observing Run. *Phys. Rev. Lett.*, 118(12):121101, 2017. [Erratum: *Phys.Rev.Lett.* 119, 029901 (2017)].

- [252] B. P. Abbott et al. Search for the isotropic stochastic background using data from Advanced LIGO’s second observing run. *Phys. Rev. D*, 100(6):061101, 2019.
- [253] R. Abbott et al. Constraints on Cosmic Strings Using Data from the Third Advanced LIGO–Virgo Observing Run. *Phys. Rev. Lett.*, 126(24):241102, 2021.
- [254] R. Abbott et al. Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo’s third observing run. *Phys. Rev. D*, 104(2):022004, 2021.
- [255] R. Abbott et al. Search for anisotropic gravitational-wave backgrounds using data from Advanced LIGO and Advanced Virgo’s first three observing runs. *Phys. Rev. D*, 104(2):022005, 2021.
- [256] Alessandro Nagar, Geraint Pratten, Gunnar Riemenschneider, and Rossella Gamba. A Multipolar Effective One Body Model for Non-Spinning Black Hole Binaries. 2019.
- [257] Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal. *Phys. Rev. D*, 93(4):044006, 2016.
- [258] Jonathan E. Thompson, Edward Fauchon-Jones, Sebastian Khan, Elisa Ni-toglia, Francesco Pannarale, Tim Dietrich, and Mark Hannam. Modeling the gravitational wave signature of neutron star black hole coalescences. *Phys. Rev. D*, 101(12):124059, 2020.
- [259] Héctor Estellés, Marta Colleoni, Cecilio García-Quirós, Sascha Husa, David Keitel, Maite Mateu-Lucena, María de Lluc Planas, and Antoni Ramos-Buades. New twists in compact binary waveform modeling: A fast time-domain model for precession. *Phys. Rev. D*, 105(8):084040, 2022.
- [260] Manuel Tiglio and Aarón Villanueva. Reduced order and surrogate models for gravitational waves. *Living Rev. Rel.*, 25(1):2, 2022.
- [261] Scott E. Field, Chad R. Galley, Frank Herrmann, Jan S. Hesthaven, Evan Ochsner, and Manuel Tiglio. Reduced basis catalogs for gravitational wave templates. *Phys. Rev. Lett.*, 106:221102, 2011.
- [262] Michael Pürrer. Frequency domain reduced order models for gravitational waves from aligned-spin compact binaries. *Class. Quant. Grav.*, 31(19):195010, 2014.

- [263] Michael Pürer. Frequency domain reduced order model of aligned-spin effective-one-body waveforms with generic mass-ratios and spins. *Phys. Rev. D*, 93(6):064041, 2016.
- [264] Scott E. Field, Chad R. Galley, Jan S. Hesthaven, Jason Kaye, and Manuel Tiglio. Fast prediction and evaluation of gravitational waveforms using surrogate models. *Phys. Rev. X*, 4(3):031006, 2014.
- [265] Roberto Cotesa, Sylvain Marsat, and Michael Pürer. Frequency domain reduced order model of aligned-spin effective-one-body waveforms with higher-order modes. *Phys. Rev. D*, 101(12):124040, 2020.
- [266] Bhooshan Gadre, Michael Pürer, Scott E. Field, Serguei Ossokine, and Vijay Varma. A fully precessing higher-mode surrogate model of effective-one-body waveforms. 3 2022.
- [267] Vijay Varma, Scott E. Field, Mark A. Scheel, Jonathan Blackman, Davide Gerosa, Leo C. Stein, Lawrence E. Kidder, and Harald P. Pfeiffer. Surrogate models for precessing binary black hole simulations with unequal masses. *Phys. Rev. Research.*, 1:033015, 2019.
- [268] Alvin J. K. Chua, Chad R. Galley, and Michele Vallisneri. Reduced-order modeling with artificial neurons for gravitational-wave inference. *Phys. Rev. Lett.*, 122(21):211101, 2019.
- [269] Chung-Hao Liao and Feng-Li Lin. Deep generative models of gravitational waveforms via conditional autoencoder. *Phys. Rev. D*, 103(12):124051, 2021.
- [270] Lucy M. Thomas, Geraint Pratten, and Patricia Schmidt. Accelerating multi-modal gravitational waveforms from precessing compact binaries with artificial neural networks. *Phys. Rev. D*, 106(10):104029, 2022.
- [271] E. A. Huerta et al. Eccentric, nonspinning, inspiral, Gaussian-process merger approximant for the detection and characterization of eccentric binary black hole mergers. *Phys. Rev. D*, 97(2):024031, 2018.
- [272] Benjamin D. Lackey, Michael Pürer, Andrea Taracchini, and Sylvain Marsat. Surrogate model for an aligned-spin effective one body waveform model of binary neutron star inspirals using Gaussian process regression. *Phys. Rev. D*, 100(2):024002, 2019.

- [273] Paraskevi Nousi, Styliani-Christina Fragkouli, Nikolaos Passalis, Panagiotis Iosif, Theocharis Apostolatos, George Pappas, Nikolaos Stergioulas, and Anastasios Tefas. Autoencoder-driven Spiral Representation Learning for Gravitational Wave Surrogate Modelling. 7 2021.
- [274] Styliani-Christina Fragkouli, Paraskevi Nousi, Nikolaos Passalis, Panagiotis Iosif, Nikolaos Stergioulas, and Anastasios Tefas. Deep Residual Error and Bag-of-Tricks Learning for Gravitational Wave Surrogate Modeling. 3 2022.
- [275] Asad Khan, E. A. Huerta, and Huihuo Zheng. Interpretable AI forecasting for numerical relativity waveforms of quasicircular, spinning, nonprecessing binary black hole mergers. *Phys. Rev. D*, 105(2):024024, 2022.
- [276] Alvin J. K. Chua, Michael L. Katz, Niels Warburton, and Scott A. Hughes. Rapid generation of fully relativistic extreme-mass-ratio-inspiral waveform templates for LISA data analysis. *Phys. Rev. Lett.*, 126(5):051102, 2021.
- [277] S. Klimenko, I. Yakushin, A. Mercer, and Guenakh Mitselmakher. Coherent method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 25:114029, 2008.
- [278] S. Klimenko, G. Vedovato, M. Drago, G. Mazzolo, G. Mitselmakher, C. Pankow, G. Prodi, V. Re, F. Salemi, and I. Yakushin. Localization of gravitational wave sources with networks of advanced detectors. *Phys. Rev. D*, 83:102001, 2011.
- [279] S. Klimenko et al. Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors. *Phys. Rev. D*, 93(4):042004, 2016.
- [280] M. Drago et al. Coherent WaveBurst, a pipeline for unmodeled gravitational-wave data analysis. 6 2020.
- [281] Adrian Macquet, Marie-Anne Bizouard, Nelson Christensen, and Michael Coughlin. Long-duration transient gravitational-wave search pipeline. *Phys. Rev. D*, 104(10):102005, 2021.
- [282] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Phys. Rev. D*, 81:062003, 2010.
- [283] J. Veitch et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library. *Phys. Rev. D*, 91(4):042003, 2015.

- [284] Gregory Ashton et al. BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy. *Astrophys. J. Suppl.*, 241(2):27, 2019.
- [285] Joshua S. Speagle. A Conceptual Introduction to Markov Chain Monte Carlo Methods. 9 2019.
- [286] John Skilling. Nested sampling for general Bayesian computation. *Bayesian Analysis*, 1(4):833–859, 2006.
- [287] Greg Ashton et al. Nested sampling for physical scientists. *Nature*, 2, 2022.
- [288] Stephen R. Green, Christine Simpson, and Jonathan Gair. Gravitational-wave parameter estimation with autoregressive neural network flows. *Phys. Rev. D*, 102(10):104057, 2020.
- [289] Michael J. Williams, John Veitch, and Chris Messenger. Nested sampling with normalizing flows for gravitational-wave inference. *Phys. Rev. D*, 103(10):103006, 2021.
- [290] Jurriaan Langendorff, Alex Kolmus, Justin Januart, and Chris Van Den Broeck. Normalizing Flows as an Avenue to Studying Overlapping Gravitational Wave Signals. *Phys. Rev. Lett.*, 130(17):171402, 2023.
- [291] Uddipta Bhardwaj, James Alvey, Benjamin Kurt Miller, Samaya Nissanke, and Christoph Weniger. Sequential simulation-based inference for gravitational wave signals. *Phys. Rev. D*, 108(4):042004, 2023.
- [292] Michael J. Williams, John Veitch, and Chris Messenger. Importance nested sampling with normalising flows. *Mach. Learn. Sci. Tech.*, 4(3):035011, 2023.
- [293] Rahul Biswas, Patrick R. Brady, Jolien D. E. Creighton, and Stephen Fairhurst. The Loudest event statistic: General formulation, properties and applications. *Class. Quant. Grav.*, 26:175009, 2009. [Erratum: Class.Quant.Grav. 30, 079502 (2013)].
- [294] Will M. Farr, Jonathan R. Gair, Ilya Mandel, and Curt Cutler. Counting and confusion: Bayesian rate estimation with multiple populations. *Phys. Rev. D*, 91:023005, Jan 2015.
- [295] Maya Fishbach, Daniel E. Holz, and Will M. Farr. Does the Black Hole Merger Rate Evolve with Redshift? *Astrophys. J. Lett.*, 863(2):L41, 2018.
- [296] Maya Fishbach, Reed Essick, and Daniel E. Holz. Does Matter Matter? Using the mass distribution to distinguish neutron stars and black holes. *Astrophys. J. Lett.*, 899:L8, 2020.

[297] Amanda M. Farah, Maya Fishbach, Reed Essick, Daniel E. Holz, and Shanika Galaudage. Bridging the Gap: Categorizing Gravitational-wave Events at the Transition between Neutron Stars and Black Holes. *Astrophys. J.*, 931(2):108, 2022.

[298] Ilya Mandel, Will M. Farr, Andrea Colonna, Simon Stevenson, Peter Tišo, and John Veitch. Model-independent inference on compact-binary observations. *Mon. Not. Roy. Astron. Soc.*, 465(3):3254–3260, 2017.

[299] Vaibhav Tiwari. VAMANA: modeling binary black hole population with minimal assumptions. *Class. Quant. Grav.*, 38(15):155007, 2021.

[300] Stefano Rinaldi and Walter Del Pozzo. (H)DPGMM: a hierarchy of Dirichlet process Gaussian mixture models for the inference of the black hole mass function. *Mon. Not. Roy. Astron. Soc.*, 509(4):5454–5466, 2021.

[301] Thomas A. Callister and Will M. Farr. Parameter-Free Tour of the Binary Black Hole Population. *Phys. Rev. X*, 14(2):021005, 2024.

[302] Anarya Ray, Ignacio Magaña Hernandez, Siddharth Mohite, Jolien Creighton, and Shasvath Kapadia. Nonparametric Inference of the Population of Compact Binaries from Gravitational-wave Observations Using Binned Gaussian Processes. *Astrophys. J.*, 957(1):37, 2023.

[303] Nicolás Yunes and Xavier Siemens. Gravitational-Wave Tests of General Relativity with Ground-Based Detectors and Pulsar Timing-Arrays. *Living Rev. Rel.*, 16:9, 2013.

[304] Abhirup Ghosh, Nathan K. Johnson-Mcdaniel, Archisman Ghosh, Chandra Kant Mishra, Parameswaran Ajith, Walter Del Pozzo, Christopher P. L. Berry, Alex B. Nielsen, and Lionel London. Testing general relativity using gravitational wave signals from the inspiral, merger and ringdown of binary black holes. *Class. Quant. Grav.*, 35(1):014002, 2018.

[305] Roberto Cotesta, Gregorio Carullo, Emanuele Berti, and Vitor Cardoso. Analysis of Ringdown Overtones in GW150914. *Phys. Rev. Lett.*, 129(11):111102, 2022.

[306] Joshua A. Faber and Frederic A. Rasio. Binary Neutron Star Mergers. *Living Rev. Rel.*, 15:8, 2012.

[307] Matthew D. Duez. Black hole-neutron star binaries. 4 2024.

- [308] Katerina Chatzioannou. Neutron star tidal deformability and equation of state constraints. *Gen. Rel. Grav.*, 52(11):109, 2020.
- [309] Ehud Nakar. Short-Hard Gamma-Ray Bursts. *Phys. Rept.*, 442:166–236, 2007.
- [310] Frédéric Piron. Gamma-Ray Bursts at high and very high energies. *Comptes Rendus Physique*, 17:617–631, 2016.
- [311] Brian D. Metzger. Kilonovae. *Living Rev. Rel.*, 23(1):1, 2020.
- [312] B. P. Abbott et al. Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *Astrophys. J. Lett.*, 848(2):L13, 2017.
- [313] A. Albert et al. Search for High-energy Neutrinos from Binary Neutron Star Merger GW170817 with ANTARES, IceCube, and the Pierre Auger Observatory. *Astrophys. J. Lett.*, 850(2):L35, 2017.
- [314] B. P. Abbott et al. Constraining the p -Mode– g -Mode Tidal Instability with GW170817. *Phys. Rev. Lett.*, 122(6):061104, 2019.
- [315] Benjamin P Abbott et al. Model comparison from LIGO–Virgo data on GW170817’s binary components and consequences for the merger remnant. *Class. Quant. Grav.*, 37(4):045006, 2020.
- [316] Simone Mastrogiiovanni, Christos Karathanasis, Jonathan Gair, Gregory Ashton, Stefano Rinaldi, Hsiang-Yu Huang, and Gergely Dálya. Cosmology with Gravitational Waves: A Review. *Annalen Phys.*, 536(2):2200180, 2024.
- [317] Rachel Gray et al. Cosmological inference using gravitational wave standard sirens: A mock data analysis. *Phys. Rev. D*, 101(12):122001, 2020.
- [318] S. Mastrogiiovanni, K. Leyde, C. Karathanasis, E. Chassande-Mottin, D. A. Steer, J. Gair, A. Ghosh, R. Gray, S. Mukherjee, and S. Rinaldi. On the importance of source population models for gravitational-wave cosmology. *Phys. Rev. D*, 104(6):062009, 2021.
- [319] Kenta Hotokezaka, Ehud Nakar, Ore Gottlieb, Samaya Nissanke, Kento Maeda, Gregg Hallinan, Kunal P. Mooley, and Adam. T. Deller. A Hubble constant measurement from superluminal motion of the jet in GW170817. *Nature Astron.*, 3(10):940–944, 2019.
- [320] M. Soares-Santos et al. First Measurement of the Hubble Constant from a Dark Standard Siren using the Dark Energy Survey Galaxies and the LIGO/Virgo Binary–Black-hole Merger GW170814. *Astrophys. J. Lett.*, 876(1):L7, 2019.

- [321] Jian-Ping Hu and Fa-Yin Wang. Hubble Tension: The Evidence of New Physics. *Universe*, 9(2):94, 2023.
- [322] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. [Erratum: *Astron. Astrophys.* 652, C4 (2021)].
- [323] Adam G. Riess et al. A Comprehensive Measurement of the Local Value of the Hubble Constant with $1 \text{ km s}^{-1} \text{ Mpc}^{-1}$ Uncertainty from the Hubble Space Telescope and the SH0ES Team. *Astrophys. J. Lett.*, 934(1):L7, 2022.
- [324] T. M. C. Abbott et al. The Dark Energy Survey: Cosmology Results With ~ 1500 New High-redshift Type Ia Supernovae Using The Full 5-year Dataset. 1 2024.
- [325] B. S. Sathyaprakash and S. V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 44:3819–3834, 1991.
- [326] J. Neyman and E. S. Pearson. On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Philosophical Transactions of the Royal Society of London Series A*, 231:289–337, January 1933.
- [327] Ian Harry, Stephen Privitera, Alejandro Bohé, and Alessandra Buonanno. Searching for Gravitational Waves from Compact Binaries with Precessing Spins. *Phys. Rev. D*, 94(2):024012, 2016.
- [328] Ian Harry, Juan Calderón Bustillo, and Alex Nitz. Searching for the full symphony of black hole binary mergers. *Phys. Rev. D*, 97(2):023004, 2018.
- [329] Yi Pan, Alessandra Buonanno, Yan-bei Chen, and Michele Vallisneri. A Physical template family for gravitational waves from precessing binaries of spinning compact objects: Application to single spin binaries. *Phys. Rev. D*, 69:104017, 2004. [Erratum: *Phys. Rev. D* 74, 029905 (2006)].
- [330] S. V. Dhurandhar and B. S. Sathyaprakash. Choice of filters for the detection of gravitational waves from coalescing binaries. 2. Detection in colored noise. *Phys. Rev. D*, 49:1707–1722, 1994.
- [331] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, 1999.
- [332] S. Babak, R. Balasubramanian, D. Churches, T. Cokelaer, and B. S. Sathyaprakash. A Template bank to search for gravitational waves from inspiralling compact binaries. I. Physical models. *Class. Quant. Grav.*, 23:5477–5504, 2006.

- [333] Thomas Cokelaer. A Template bank to search for gravitational waves from inspiralling compact binaries. II. Phenomenological model. *Class. Quant. Grav.*, 24:6227–6242, 2007.
- [334] T. A. Apostolatos. Search templates for gravitational waves from precessing, inspiraling binaries. *Phys. Rev. D*, 52:605–620, 1995.
- [335] Reinhard Prix. Template-based searches for gravitational waves: Efficient lattice covering of flat parameter spaces. *Class. Quant. Grav.*, 24:S481–S490, 2007.
- [336] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, 1996.
- [337] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. Gravitational waves from coalescing binaries: Detection strategies and Monte Carlo estimation of parameters. *Phys. Rev. D*, 53:3033–3055, 1996. [Erratum: Phys.Rev.D 54, 1860 (1996)].
- [338] Duncan A. Brown, Ian Harry, Andrew Lundgren, and Alexander H. Nitz. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Phys. Rev. D*, 86:084017, 2012.
- [339] Thomas Cokelaer. Gravitational waves from inspiralling compact binaries: Hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev. D*, 76:102004, 2007.
- [340] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg. Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries. *Phys. Rev. D*, 77:104017, May 2008.
- [341] Soumen Roy, Anand S. Sengupta, and Parameswaran Ajith. Effectual template banks for upcoming compact binary searches in Advanced-LIGO and Virgo data. *Phys. Rev. D*, 99(2):024048, 2019.
- [342] Javier Roulet, Liang Dai, Tejaswi Venumadhav, Barak Zackay, and Matias Zaldarriaga. Template Bank for Compact Binary Coalescence Searches in Gravitational Wave Data: A General Geometric Placement Algorithm. *Phys. Rev. D*, 99(12):123022, 2019.
- [343] Chad Hanna et al. Binary tree approach to template placement for searches for gravitational waves from compact binary mergers. *Phys. Rev. D*, 108(4):042003, 2023.

- [344] Bruce Allen. Optimal template banks. *Phys. Rev. D*, 104(4):042005, 2021.
- [345] C. Messenger, R. Prix, and M. A. Papa. Random template banks and relaxed lattice coverings. *Phys. Rev. D*, 79:104017, 2009.
- [346] Adam Coogan, Thomas D. P. Edwards, Horng Sheng Chia, Richard N. George, Katherine Freese, Cody Messick, Christian N. Setzer, Christoph Weniger, and Aaron Zimmerman. Efficient gravitational wave template bank generation with differentiable waveforms. *Phys. Rev. D*, 106(12):122001, 2022.
- [347] Bruce Allen. Performance of random template banks. *Phys. Rev. D*, 105(10):102003, 2022.
- [348] Stanislav Babak. Building a stochastic template bank for detecting massive black hole binaries. *Class. Quant. Grav.*, 25:195011, 2008.
- [349] Ian W. Harry, Bruce Allen, and B. S. Sathyaprakash. A Stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev. D*, 80:104014, 2009.
- [350] Gian Mario Manca and Michele Vallisneri. Cover art: Issues in the metric-guided and metric-less placement of random and stochastic template banks. *Phys. Rev. D*, 81:024004, 2010.
- [351] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, and A. J. Weinstein. Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins. *Phys. Rev. D*, 89(8):084041, 2014.
- [352] Tito Dal Canton and Ian W. Harry. Designing a template bank to observe compact binary coalescences in Advanced LIGO’s second observing run. 5 2017.
- [353] Debnandini Mukherjee et al. Template bank for spinning compact binary mergers in the second observation run of Advanced LIGO and the first observation run of Advanced Virgo. *Phys. Rev. D*, 103(8):084047, 2021.
- [354] Ananya Bandopadhyay, Brendan Reed, Surendra Padamata, Erick Leon, C. J. Horowitz, Duncan A. Brown, David Radice, F. J. Fattoyev, and J. Piekarewicz. Detectability of subsolar mass neutron stars through a template bank search. *Phys. Rev. D*, 107(10):103012, 2023.
- [355] Tito Dal Canton et al. Implementing a search for aligned-spin neutron star-black hole systems with advanced ground based gravitational wave detectors. *Phys. Rev. D*, 90(8):082004, 2014.

- [356] Samantha A. Usman et al. The PyCBC search for gravitational waves from compact binary coalescence. *Class. Quant. Grav.*, 33(21):215004, 2016.
- [357] Alexander H. Nitz, Thomas Dent, Tito Dal Canton, Stephen Fairhurst, and Duncan A. Brown. Detecting binary compact-object mergers with gravitational waves: Understanding and Improving the sensitivity of the PyCBC search. *Astrophys. J.*, 849(2):118, 2017.
- [358] Gareth S. Davies, Thomas Dent, Márton Tápai, Ian Harry, Connor McIsaac, and Alexander H. Nitz. Extending the PyCBC search for gravitational waves from compact binary mergers to a global network. *Phys. Rev. D*, 102(2):022004, 2020.
- [359] Connor R. Weaving, Laura K. Nuttall, Ian W. Harry, Shichao Wu, and Alexander Nitz. Adapting the PyCBC pipeline to find and infer the properties of gravitational waves from massive black hole binaries in LISA. *Class. Quant. Grav.*, 41(2):025006, 2024.
- [360] T. Adams, D. Buskulic, V. Germain, G. M. Guidi, F. Marion, M. Montani, B. Mours, F. Piergiovanni, and G. Wang. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *Class. Quant. Grav.*, 33(17):175012, 2016.
- [361] F. Aubin et al. The MBTA pipeline for detecting compact binary coalescences in the third LIGO–Virgo observing run. *Class. Quant. Grav.*, 38(9):095004, 2021.
- [362] Jing Luan, Shaun Hooper, Linqing Wen, and Yanbei Chen. Towards low-latency real-time detection of gravitational waves from compact binary coalescences in the era of advanced detectors. *Phys. Rev. D*, 85:102002, 2012.
- [363] Qi Chu. *Low-latency detection and localization of gravitational waves from compact binary coalescences*. PhD thesis, Western Australia U., 2017.
- [364] Qi Chu et al. SPIIR online coherent pipeline to search for gravitational waves from compact binary coalescences. *Phys. Rev. D*, 105(2):024023, 2022.
- [365] Collin Capano. *Searching for Gravitational Waves from Compact Binary Coalescence Using LIGO and Virgo Data*. PhD thesis, Syracuse U., 2011.
- [366] S. Babak et al. Searching for gravitational waves from binary coalescence. *Phys. Rev. D*, 87(2):024033, 2013.
- [367] Stephen P. Vitetta, Satyanarayan R. P. Mohapatra, Parameswaran Ajith, Kipp Cannon, Nickolas Fotopoulos, Melissa A. Frei, Chad Hanna, Alan J. Weinstein,

and John T. Whelan. Improving the sensitivity of a search for coalescing binary black holes with nonprecessing spins in gravitational wave data. *Phys. Rev. D*, 89(2):024003, 2014.

[368] Collin Capano, Ian Harry, Stephen Privitera, and Alessandra Buonanno. Implementing a search for gravitational waves from binary black holes with non-precessing spin. *Phys. Rev. D*, 93(12):124007, 2016.

[369] Tejaswi Venumadhav, Barak Zackay, Javier Roulet, Liang Dai, and Matias Zaldarriaga. New search pipeline for compact binary mergers: Results for binary black holes in the first observing run of Advanced LIGO. *Phys. Rev. D*, 100(2):023011, 2019.

[370] B. P. Abbott et al. Search for intermediate mass black hole binaries in the first and second observing runs of the Advanced LIGO and Virgo network. *Phys. Rev. D*, 100(6):064064, 2019.

[371] Rich Abbott et al. Search for intermediate-mass black hole binaries in the third observing run of Advanced LIGO and Advanced Virgo. *Astron. Astrophys.*, 659:A84, 2022.

[372] Koustav Chandra, V. Villa-Ortega, T. Dent, C. McIsaac, Archana Pai, I. W. Harry, G. S. Cabourn Davies, and K. Soni. An optimized PyCBC search for gravitational waves from intermediate-mass black hole mergers. *Phys. Rev. D*, 104:042004, 2021.

[373] Koustav Chandra, Juan Calderón Bustillo, Archana Pai, and I. W. Harry. First gravitational-wave search for intermediate-mass black hole mergers with higher-order harmonics. *Phys. Rev. D*, 106(12):123003, 2022.

[374] Alexander Harvey Nitz and Yi-Fan Wang. Search for Gravitational Waves from High-Mass-Ratio Compact-Binary Mergers of Stellar Mass and Subsolar Mass Black Holes. *Phys. Rev. Lett.*, 126(2):021103, 2021.

[375] Alexander H. Nitz and Yi-Fan Wang. Search for gravitational waves from the coalescence of sub-solar mass and eccentric compact binaries. 2 2021.

[376] R. Abbott et al. Search for Subsolar-Mass Binaries in the First Half of Advanced LIGO's and Advanced Virgo's Third Observing Run. *Phys. Rev. Lett.*, 129(6):061104, 2022.

[377] Alexander H. Nitz and Yi-Fan Wang. Search for Gravitational Waves from the Coalescence of Subsolar-Mass Binaries in the First Half of Advanced LIGO and Virgo's Third Observing Run. *Phys. Rev. Lett.*, 127(15):151101, 2021.

- [378] R. Abbott et al. Search for subsolar-mass black hole binaries in the second part of Advanced LIGO’s and Advanced Virgo’s third observing run. *Mon. Not. Roy. Astron. Soc.*, 524(4):5984–5992, 2023. [Erratum: Mon. Not. Roy. Astron. Soc. 526, 6234 (2023)].
- [379] Kipp Cannon et al. Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence. *Astrophys. J.*, 748:136, 2012.
- [380] Shio Sakon et al. Template bank for compact binary mergers in the fourth observing run of Advanced LIGO, Advanced Virgo, and KAGRA. *Phys. Rev. D*, 109(4):044066, 2024.
- [381] Soichiro Morisaki and Vivien Raymond. Rapid Parameter Estimation of Gravitational Waves from Binary Neutron Star Coalescence using Focused Reduced Order Quadrature. *Phys. Rev. D*, 102(10):104020, 2020.
- [382] Bruce Allen. χ^2 time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, 2005.
- [383] Peter Shawhan and Evan Ochsner. A New waveform consistency test for gravitational wave inspiral searches. *Class. Quant. Grav.*, 21:S1757–S1766, 2004.
- [384] Alexander Harvey Nitz. Distinguishing short duration noise transients in LIGO data to improve the PyCBC search for gravitational waves from high mass binary black hole mergers. *Class. Quant. Grav.*, 35(3):035016, 2018.
- [385] Sanjeev Dhurandhar, Anuradha Gupta, Bhooshan Gadre, and Sukanta Bose. A unified approach to χ^2 discriminators for searches of gravitational waves from compact binary coalescences. *Phys. Rev. D*, 96(10):103018, 2017.
- [386] V. Gayathri, P. Bacon, A. Pai, E. Chassande-Mottin, F. Salemi, and G. Vedovato. Astrophysical signal consistency test adapted for gravitational-wave transient searches. *Phys. Rev. D*, 100(12):124022, 2019.
- [387] Connor McIsaac and Ian Harry. Using machine learning to autotune chi-squared tests for gravitational wave searches. *Phys. Rev. D*, 105(10):104056, 2022.
- [388] Kipp Cannon, Chad Hanna, and Drew Keppel. Method to estimate the significance of coincident gravitational-wave observations from compact binary coalescence. *Phys. Rev. D*, 88(2):024025, 2013.

[389] Kipp Cannon, Chad Hanna, and Jacob Peoples. Likelihood-Ratio Ranking Statistic for Compact Binary Coalescence Candidates with Rate Estimation. 4 2015.

[390] Reed Essick, Patrick Godwin, Chad Hanna, Lindy Blackburn, and Erik Kat-savounidis. iDQ: Statistical Inference of Non-Gaussian Noise with Auxiliary Degrees of Freedom in Gravitational-Wave Detectors. 5 2020.

[391] Patrick Godwin et al. Incorporation of Statistical Data Quality Information into the GstLAL Search Analysis. 10 2020.

[392] B. P. Abbott et al. The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. *Astrophys. J. Lett.*, 833(1):L1, 2016.

[393] B. P. Abbott et al. Supplement: The Rate of Binary Black Hole Mergers Inferred from Advanced LIGO Observations Surrounding GW150914. *Astrophys. J. Suppl.*, 227(2):14, 2016.

[394] Thomas J. Loredo and Ira M. Wasserman. Inferring the Spatial and Energy Distribution of Gamma-Ray Burst Sources. I. Methodology. , 96:261, January 1995.

[395] F. Guglielmetti, R. Fischer, and V. Dose. Background-Source Separation in astronomical images with Bayesian probability theory (I): the method. *Mon. Not. Roy. Astron. Soc.*, 396:165, 2009.

[396] Shasvath J. Kapadia et al. A self-consistent method to estimate the rate of compact binary coalescences with a Poisson mixture model. *Class. Quant. Grav.*, 37(4):045007, 2020.

[397] B. P. Abbott et al. Search for Subsolar Mass Ultracompact Binaries in Advanced LIGO's Second Observing Run. *Phys. Rev. Lett.*, 123(16):161102, 2019.

[398] Alexander H. Nitz and Yi-Fan Wang. Broad search for gravitational waves from subsolar-mass binaries through LIGO and Virgo's third observing run. *Phys. Rev. D*, 106(2):023024, 2022.

[399] Antoni Ramos-Buades, Shubhanshu Tiwari, Maria Haney, and Sascha Husa. Impact of eccentricity on the gravitational-wave searches for binary black holes: High mass case. *Phys. Rev. D*, 102:043005, Aug 2020.

- [400] Belinda D. Cheeseboro and Paul T. Baker. Method for detecting highly eccentric binaries with a gravitational wave burst search. *Phys. Rev. D*, 104:104016, Nov 2021.
- [401] Alexander H. Nitz, Amber Lennon, and Duncan A. Brown. Search for Eccentric Binary Neutron Star Mergers in the first and second observing runs of Advanced LIGO. *Astrophys. J.*, 890:1, 12 2019.
- [402] B. P. Abbott et al. Search for Eccentric Binary Black Hole Mergers with Advanced LIGO and Advanced Virgo during their First and Second Observing Runs. *Astrophys. J.*, 883(2):149, 2019.
- [403] Yi-Fan Wang and Alexander H. Nitz. Prospects for detecting gravitational waves from eccentric subsolar mass compact binaries. *Astrophys. J.*, 912(1):53, 2021.
- [404] Ian W. Harry, Alexander H. Nitz, Duncan A. Brown, Andrew P. Lundgren, Evan Ochsner, and Drew Keppel. Investigating the effect of precession on searches for neutron-star-black-hole binaries with Advanced LIGO. *Phys. Rev. D*, 89(2):024010, 2014.
- [405] Nathaniel Indik, K. Haris, Tito Dal Canton, Henning Fehrmann, Badri Krishnan, Andrew Lundgren, Alex B. Nielsen, and Archana Pai. Stochastic template bank for gravitational wave searches for precessing neutron-star–black-hole coalescence events. *Phys. Rev. D*, 95(6):064056, 2017.
- [406] Connor McIsaac, Charlie Hoy, and Ian Harry. Search technique to observe precessing compact binary mergers in the advanced detector era. *Phys. Rev. D*, 108(12):123016, 2023.
- [407] Juan Calderón Bustillo, Sascha Husa, Alicia M. Sintes, and Michael Pürrer. Impact of gravitational radiation higher order modes on single aligned-spin gravitational wave searches for binary black holes. *Phys. Rev. D*, 93(8):084019, 2016.
- [408] Cameron Mills and Stephen Fairhurst. Measuring gravitational-wave higher-order multipoles. *Phys. Rev. D*, 103(2):024042, 2021.
- [409] Digvijay Wadekar, Tejaswi Venumadhav, Ajit Kumar Mehta, Javier Roulet, Seth Olsen, Jonathan Mushkin, Barak Zackay, and Matias Zaldarriaga. A new approach to template banks of gravitational waves with higher harmonics: reducing matched-filtering cost by over an order of magnitude. 10 2023.
- [410] Stefano Schmidt. `mbank` - metric bank generation for gravitational waves data analysis. <https://mbank.readthedocs.io/en/latest/>.

- [411] Sebastian Khan and Rhys Green. Gravitational-wave surrogate models powered by artificial neural networks. *Phys. Rev. D*, 103(6):064015, 2021.
- [412] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. 2019.
- [413] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios. *Neural spline flows*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [414] Ivan Kobyzev, Simon J.D. Prince, and Marcus A. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, nov 2021.
- [415] George Papamakarios. Neural density estimation and likelihood-free inference, 2019.
- [416] Mathieu Germain, Karol Gregor, Iain Murray, and Hugo Larochelle. MADE: masked autoencoder for distribution estimation. *CoRR*, abs/1502.03509, 2015.
- [417] George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. 2017.
- [418] Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron C. Courville. Neural autoregressive flows. *CoRR*, abs/1804.00779, 2018.
- [419] S. Kullback and R. A. Leibler. On Information and Sufficiency. *The Annals of Mathematical Statistics*, 22(1):79 – 86, 1951.
- [420] Kevin P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [421] LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite. free software (GPL), 2018.
- [422] Cecilio García-Quirós, Marta Colleoni, Sascha Husa, Héctor Estellés, Geraint Pratten, Antoni Ramos-Buades, Maite Mateu-Lucena, and Rafel Jaume. Multimode frequency-domain model for the gravitational wave signal from non-precessing black-hole binaries. *Phys. Rev. D*, 102(6):064002, 2020.
- [423] R. Abbott et al. Advanced ligo anticipated sensitivity curves. <https://dcc.ligo.org/LIGO-T0900288/public>.

- [424] Xingjiang Zhu, Eric Thrane, Stefan Oslowski, Yuri Levin, and Paul D. Lasky. Inferring the population properties of binary neutron stars with gravitational-wave measurements of spin. *Phys. Rev. D*, 98:043002, 2018.
- [425] R. Abbott et al. Noise curves used for simulations in the update of the observing scenarios paper. <https://dcc.ligo.org/LIGO-T2000012/public>.
- [426] S. Babak, H. Grote, M. Hewitson, H. Luck, and K. A. Strain. Signal based vetoes for the detection of gravitational waves from inspiralling compact binaries. *Phys. Rev. D*, 72:022002, 2005.
- [427] 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):102003, 2014.
- [428] Timothy D. Gebhard, Niki Kilbertus, Ian Harry, and Bernhard Schölkopf. Convolutional neural networks: a magic bullet for gravitational-wave detection? *Phys. Rev. D*, 100(6):063015, 2019.
- [429] Marlin B. Schäfer, Frank Ohme, and Alexander H. Nitz. Detection of gravitational-wave signals from binary neutron star mergers using machine learning. *Phys. Rev. D*, 102(6):063015, 2020.
- [430] Marlin B. Schäfer, Ondřej Zelenka, Alexander H. Nitz, Frank Ohme, and Bernd Brügmann. Training strategies for deep learning gravitational-wave searches. *Phys. Rev. D*, 105(4):043002, 2022.
- [431] Gregory Baltus, Justin Janquart, Melissa Lopez, Amit Reza, Sarah Caudill, and Jean-Rene Cudell. Convolutional neural networks for the detection of the early inspiral of a gravitational-wave signal. *Phys. Rev. D*, 103:102003, 2021.
- [432] João D. Álvares, José A. Font, Felipe F. Freitas, Osvaldo G. Freitas, António P. Morais, Solange Nunes, Antonio Onofre, and Alejandro Torres-Forné. Gravitational-wave parameter inference using Deep Learning. 11 2020.
- [433] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *CoRR*, abs/1410.8516, 2014.
- [434] Laurent Dinh, Jascha Narain Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *ArXiv*, abs/1605.08803, 2016.
- [435] Yoshua Bengio and Samy Bengio. Modeling high-dimensional discrete data with multi-layer neural networks. In S. Solla, T. Leen, and K. Müller, editors, *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.

Bibliography

- [436] Jens Behrmann, David Kristjanson Duvenaud, and Jörn-Henrik Jacobsen. Invertible residual networks. In *International Conference on Machine Learning*, 2018.
- [437] Nathaniel Indik, Henning Fehrmann, Franz Harke, Badri Krishnan, and Alex B. Nielsen. Reducing the number of templates for aligned-spin compact binary coalescence gravitational wave searches using metric-agnostic template nudging. *Phys. Rev. D*, 97(12):124008, 2018.
- [438] Ian W. Harry, Alexander H. Nitz, Duncan A. Brown, Andrew P. Lundgren, Evan Ochsner, and Drew Keppel. Investigating the effect of precession on searches for neutron-star-black-hole binaries with advanced ligo. *Phys. Rev. D*, 89:024010, Jan 2014.
- [439] Abbott, R. and others. Noise curves used for Simulations in the update of the Observing Scenarios Paper. <https://dcc.ligo.org/LIGO-T2000012/public>, 2022. [Online; accessed 29-January-2024].
- [440] Theocharis A. Apostolatos. Construction of a template family for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 54:2421–2437, 1996.
- [441] Alessandra Buonanno, Yanbei Chen, Yi Pan, Hideyuki Tagoshi, and Michele Vallisneri. Detecting gravitational waves from precessing binaries of spinning compact objects. II. Search implementation for low-mass binaries. *Phys. Rev. D*, 72:084027, 2005.
- [442] Digvijay Wadekar, Javier Roulet, Tejaswi Venumadhav, Ajit Kumar Mehta, Barak Zackay, Jonathan Mushkin, Seth Olsen, and Matias Zaldarriaga. New black hole mergers in the LIGO-Virgo O3 data from a gravitational wave search including higher-order harmonics. 12 2023.
- [443] Digvijay Wadekar, Tejaswi Venumadhav, Javier Roulet, Ajit Kumar Mehta, Barak Zackay, Jonathan Mushkin, and Matias Zaldarriaga. A new search pipeline for gravitational waves with higher-order modes using mode-by-mode filtering. 5 2024.
- [444] Kipp Cannon, Adrian Chapman, Chad Hanna, Drew Keppel, Antony C. Searle, and Alan J. Weinstein. Singular value decomposition applied to compact binary coalescence gravitational-wave signals. *Phys. Rev. D*, 82:044025, 2010.
- [445] Michele Vallisneri, Jonah Kanner, Roy Williams, Alan Weinstein, and Branson Stephens. The LIGO Open Science Center. *J. Phys. Conf. Ser.*, 610(1):012021, 2015.

- [446] Vaibhav Tiwari. Estimation of the Sensitive Volume for Gravitational-wave Source Populations Using Weighted Monte Carlo Integration. *Class. Quant. Grav.*, 35(14):145009, 2018.
- [447] David W. Hogg. Distance measures in cosmology. 5 1999.
- [448] Ethan Payne, Sophie Hourihane, Jacob Golomb, Rhiannon Udall, Richard Udall, Derek Davis, and Katerina Chatzioannou. Curious case of GW200129: Interplay between spin-precession inference and data-quality issues. *Phys. Rev. D*, 106(10):104017, 2022.
- [449] Ronaldas Macas, Andrew Lundgren, and Gregory Ashton. Revisiting GW200129 with machine learning noise mitigation: it is (still) precessing. 11 2023.
- [450] M. Tse et al. Quantum-Enhanced Advanced LIGO Detectors in the Era of Gravitational-Wave Astronomy. *Phys. Rev. Lett.*, 123(23):231107, 2019.

