



Fast analysis of long gravitational wave signals

JACOPO TISSINO
ADVISOR: JAN HARMS

Gran Sasso Science Institute
Laboratori Nazionali del Gran Sasso

2026-06-26

Table of contents

1. Introduction	4
1.1. Personal contributions	5
1.2. Acknowledgements	6
I. Parameter estimation and model selection	7
2. Statistical methods	8
2.1. Inference and model comparison	8
2.1.1. Posterior samples	9
2.1.2. p - p tests	14
2.1.3. p - p test failure modes	17
2.1.4. Model comparison	20
2.2. Nested sampling	23
2.2.1. Prior compression	24
2.2.2. Computing evidence and posterior	26
2.2.3. Sampling options	28
2.2.4. Model comparison with nested sampling	30
2.2.5. Commonly used plots	31
2.2.6. Information gain	36
2.2.7. Acceleration	39
2.3. Simulation-based inference	41
2.3.1. Normalizing flows	41
2.3.2. Neural posterior estimation	42
2.3.3. Importance sampling	43
2.4. Fisher matrix	47
2.4.1. Analytic Fisher calculations	49
3. Gravitational wave parameter estimation	54
3.1. Fourier transform basics	54
3.2. The Whittle likelihood	55
3.2.1. Non-periodic data: windowing	57
3.2.2. Non-periodic data: time domain parameter estimation	59

3.2.3.	Strain plotting and whitening	61
3.3.	Matched filtering and the signal-to-noise ratio	63
3.3.1.	SNR thresholds	64
3.3.2.	Noise plotting and characteristic strain	65
3.4.	Likelihood acceleration	66
3.4.1.	Relative binning	67
3.4.2.	Reduced order quadrature	69
3.5.	Marginalization	70
3.5.1.	Reference phase	70
II.	Anatomy of a gravitational wave signal	72
4.	Tests of General Relativity with GW250114	73
4.1.	The signal	73
4.2.	Ringdown	75
4.3.	Detection of the 221 overtone	76
4.3.1.	Spectroscopic test of the Kerr metric	78
4.4.	Area law	80
5.	Intrinsic parameters and waveform modeling	84
5.1.	Component masses	84
5.1.1.	The component masses of GW250114	85
5.1.2.	At lower frequencies	86
5.1.3.	Prior distribution	87
5.2.	Spin and orbital angular momentum	87
5.2.1.	Prior distribution	90
5.3.	Eccentricity	90
5.3.1.	At lower frequencies	93
5.4.	Tidal parameters	96
5.5.	Waveform modelling	98
5.5.1.	Post-Newtonian	99
5.5.2.	Effective One Body	99
5.5.3.	Surrogate models: the mlgw_bns example	102
6.	Extrinsic parameters and signal projection	105
6.1.	Extrinsic parameters	105
6.1.1.	Phase and inclination angle	105
6.1.2.	Sky position and polarization	109
6.1.3.	Distance and redshift	111
6.1.4.	Timing	113

6.2. Stationary phase approximation	115
6.2.1. SPA as linearization in frequency	116
6.3. Calibration	117
6.4. Antenna patterns	120
6.4.1. Interferometer antenna patterns	121
6.4.2. Antenna patterns for the Lunar Gravitational Wave Antenna	121
III. Future gravitational wave detectors	124
7. Einstein Telescope	125
7.1. Detector concept	125
7.2. Data analysis challenges	127
7.2.1. GW250114 as seen by a triangular Einstein Telescope	128
7.3. Neural posterior estimation for high-mass black holes	131
7.3.1. Δ shape	136
7.3.2. 2L shape	139
7.3.3. Statistical validation	139
8. Lunar Gravitational Wave Antenna	143
8.1. Detector concept	143
8.1.1. Site selection	145
8.1.2. Measurement principle and displacement noise	147
8.1.3. The lunar seismic response	152
8.1.4. Projected gravitational wave strain	153
8.2. Scientific prospects	161
8.2.1. Lunar science	161
8.2.2. Gravitational wave sources	162
8.2.3. Parameter estimation for stellar-mass binaries	168
8.2.4. Impact of non-zero Gaussian noise	174
8.3. Reference frames for long observations	183
8.3.1. Shifting reference frame in post-processing	186
8.3.2. Spanned area scaling validation	190
8.3.3. Effect on sampling	198
9. Conclusions	200
References	202

1. Introduction

We offer the reader a tour of the gravitational wave science enabled by the study of compact binary coalescences. This field is wide and multi-faceted: we center our discussion on data analysis, which encompasses the statistical and computational methods used to infer the properties of the astrophysical objects whose emission we measure.

The relative simplicity of the data produced by gravitational wave detectors allows us to employ statistically robust Bayesian methods. Starting from a detailed discussion of these, we explore how they can help shed light on the properties of compact objects. We start from current gravitational wave detectors and then move to the ones whose construction is planned: each of these brings its own scientific potential, as well as new challenges.

Chapter 2 introduces the general tools required for inference; Chapter 3 illustrates how they can be applied to the analysis of timeseries. In both, we outline what approaches can improve computational efficiency, as well as how the robustness of our methods can be assessed.

The next three chapters summarize the main ways the data measured at the detectors can vary depending on the compact binary system that emitted it. In Chapter 4 we introduce the GW250114 signal, which as of this writing holds the record as the clearest-ever gravitational wave detection: we use it as our reference example wherever possible in later chapters. In Chapter 5 we outline how the signal produced by a compact binary merger is characterized, while in Chapter 6 we take it from the source to the detector, by discussing the geometric relation between us and the emitter.

Finally, in Chapter 7 and Chapter 8 we look toward the next decades with two concepts for future detectors. Their precise data in new frequency bands bring exciting scientific prospects, but they also come with high requirements for our data analysis infrastructure: the current approaches to the problem are not trivially extended, and advancements will be needed to ensure that we can draw robust scientific conclusions for a moderate cost in time and energy.

This thesis is simultaneously compiled with [quarto](#) into both HTML — as a static website — and PDF. The HTML version is interactive, with hyperlinks, the possibility to see internal and external references by mouse over, and a dynamic table of contents. Also, several figures are rendered from python code, which can be seen directly within the website,

but is omitted from the PDF. This literate programming-like approach [105] allows for reproducibility and transparency. Its source code is publically available [here](#) (to be made public upon graduation).

1.1. Personal contributions

This thesis was deliberately written weaving established results and novel, personal contributions in order to improve its logical flow and, hopefully, its readability. This operation may obfuscate the distinction: this section addresses the problem by explicitly outlining it. For clarity, this is written in the first person; while in the rest of the thesis, the pronoun “we” is used. I opted for it because all the results presented were obtained with at least some degree of collaboration, even for the papers I led.

I am part of several collaborations: LIGO-Virgo-KAGRA, Einstein Telescope and Lunar Gravitational Wave Antenna. All three of these have recently published major works, which are formally authored by the whole collaboration.

The LIGO-Virgo-KAGRA collaboration recently concluded its fourth observing run, whose data is currently being analyzed. I was part of the team which wrote the discovery paper for the exceptional gravitational wave event GW250114 (Chapter 4) [2]. I authored all figures in the main text except for figure 2, and co-authored figures 2 and 3 in the Supplement. I was also responsible for the data release [113], as well as coordinating with the outreach team for the creation of various items, such as the infographic and image shown [here](#). The analysis of this event helped in identifying an issue (Section 3.2.1) in the way parameter estimation was performed in the collaboration, and I contributed to the presentation and correctness of the paper which outlined its fix [173]. I also contributed to the collaboration in other ways, such as acting as an internal paper and waveform model reviewee. By writing an interface for the eccentric waveform model TEOBResumS-Dalí, I contributed to a reanalysis of the first gravitational wave event performed including both eccentricity and spin-induced orbital precession [81].

The Einstein Telescope collaboration (Chapter 7) published a white paper [3] outlining the detector’s capabilities. I contributed to its section 8, with a discussion of the computational bottlenecks related to waveform evaluation. These are reported in this thesis’ Section 7.2, which also includes the results of a series of papers led by Filippo Santoliquido [157, 159, 158], related to population-scale data analysis with the Einstein Telescope. We collaborated closely, and besides editing all the text, I work on the theoretical aspects for many of them, like the theoretical Bayesian framework for the classification of binaries [157] or the understanding of the arrival time bimodality with ET- Δ [159].

The Lunar Gravitational Wave Antenna (LGWA) collaboration also published a white paper [6]. I contributed to sections of it relating to sky localization capabilities and the detection of binaries containing white dwarfs. The problem of performing parameter estimation with the LGWA has been a long-standing one. I included some preliminary results on it both in the white paper [6] and in a work we published comparing different concepts for lunar gravitational wave detection [56]; in the latter, I also contributed to the analysis of some concepts other than the seismic one, such as the effect of dust on an unshielded laser-interferometric detector. I led a paper on the generic solution to the LGWA parameter estimation problem [183], where we show that it must generically be performed in a coordinate frame comoving with the Solar System Barycenter, which should be centered at a location close to the detector for better computational efficiency. I was also involved in a recent work on LGWA parameter estimation for a massive black hole binary [96]: I performed cross-checks to ensure that the analysis was consistent; indeed, the signal lasted for a short enough time in band not to require an accurate treatment of the lunar orbit around the Sun.

I contributed to the development and maintenance of the software package GWFish [68], and in the validation of the parameter estimation forecasts it produced against GWTC-3 analyses [67]. I helped in constructing a procedure to extract information from summary data from the LHAASO collaboration [16].

Finally, this thesis includes (Section 5.5.3) the results from a paper I led [182] related to surrogate modelling for the gravitational waveforms emitted by neutron star binaries.

1.2. Acknowledgements

While writing this thesis I benefited from the help of many mentors, colleagues and friends.

I would like to thank my supervisor, Jan Harms, as well as the referees who read through this thesis to ensure it was in the best shape possible: Enrico Barausse, Giovanni Benato, Isobel Romero-Shaw.

Several friends and colleagues also helped me with the thesis: I would specifically like to thank Pierpaolo Bilotto, Gregorio Carullo, Rossella Gamba, Sara Gliorio, Vishnu Sanjay, Filippo Santoliquido, Michael Williams.

Many more people supported me through these years, with a list that would be too long to make comprehensive. I would like to thank friends and family from the many communities I've had the honor of joining, in L'Aquila, Pordenone, Udine, Jena; as well as the many colleagues I have worked with, scattered throughout Europe and other continents.

Part I.

Parameter estimation and model selection

2. Statistical methods

In this chapter, we review of the theory behind Bayesian parameter estimation, including posterior estimation, model comparison, as well as a detailed description of a commonly used diagnostic tool: the p - p plot (Section 2.1.2)

Then, we provide a description of the main statistical methods used to perform state-of-the-art parameter estimation for gravitational waves, including Nested Sampling (Section 2.2) and Neural Posterior Estimation (Section 2.3.2).

2.1. Inference and model comparison

This section contains the basic aspects of Bayesian parameter estimation, the backbone of all gravitational wave data analysis. This theory is of general validity, and it does not require us to restrict ourselves to the gravitational wave application quite yet.

When performing the analysis of some data d according to some model M parameterized by a finite number of parameters θ , we employ Bayes' theorem, which is derived by expressing the joint probability $\text{prob}(d, \theta|M)$ in two different ways, making use of the identity $\text{prob}(a, b) = \text{prob}(a|b)\text{prob}(b)$:

$$\text{prob}(d, \theta|M) = \mathcal{L}(d|\theta, M)\pi(\theta|M) = p(\theta|d, M)\mathcal{Z}(d|M). \quad (2.1)$$

The entries in this equation are:

- the *likelihood* $\mathcal{L}$ of the data given the model and a specific choice of its parameters — it is a probability distribution over the data (i.e. $\int \mathcal{L}(d|\theta, M)dd = 1$), but we interpret it as a function of the parameters θ , since the data we observed are fixed;
- the *prior* π — a representation of our beliefs about the model's parameters before seeing the data, it is a probability distribution over the parameters;
- the *posterior* p for the model's parameters — a representation of our updated beliefs about them after seeing the data, it is a probability distribution over the parameters;
- the *evidence* $\mathcal{Z}$ for the model — a normalization constant for the posterior, or the likelihood for the data marginalized over all the model's parameters.

An explicit formulation of the likelihood and prior is generally part of the problem definition, and can often be readily computed at any given value of the parameters.¹

If the parameter space is high-dimensional, evaluating the posterior on a grid quickly becomes unpractical. A common solution to this issue is to use stochastic methods, which allow us to obtain a set of samples θ_i distributed according to the posterior and from which we can obtain information as discussed in Section 2.1.1.

There exist many different such algorithms. Several of them can probe the posterior distribution $p(\theta|d, M) \propto \mathcal{L}(d|\theta, M)\pi(\theta|M)$ without computing its normalization constant $\mathcal{Z}$, the evidence: for example, the Metropolis-Hastings algorithm, a flavor of Markov-Chain Monte Carlo [94], operates by accepting or rejecting proposals based on the ratio of their probabilities, which is normalization-independent. Nested Sampling, which we shall discuss in more detail in Section 2.2, was developed as a method for computing the evidence, but it also yields the posterior as a by-product. Machine learning-based methods such as Simulation-Based Inference (Section 2.3) can perform this task even faster.

2.1.1. Posterior samples

What we refer to as “computing the posterior” means obtaining a useful representation for it. For example, we are often interested in the *marginal distributions* of single parameters. The marginal posterior for the parameter θ_i is defined by integrating over all the other parameters, *i.e.* *marginalizing*:

$$p(\theta_i|d) = \int \left(\prod_{k \neq i} d\theta_k \right) p(\theta|d).$$

The marginal distribution over two or more parameters can be defined analogously, by integrating along all dimensions except for the ones at hand. Marginal distributions are the best approximation for what our posterior belief about the given parameter(s) should be; they include all the uncertainty on that parameter driven by the uncertainty in other, related parameters. Figure 2.1 illustrates this by comparing the marginal distribution to the *conditional* distribution, obtained by fixing the values of the other parameters to some value:

$$p(\theta_i|d, \{\theta_k^0\}_{k \neq i}) = p(\theta|d) \text{ computed at } \{\theta_k = \theta_k^0\}_{k \neq i}.$$

Conditional distributions typically have less variance than their respective marginals, but unless we have a reason (such as a different observation) to specifically choose a given value for the other parameters, they are inaccessible in real parameter estimation.

¹This is not true for all approaches: in Simulation-Based Inference (SBI, see Section 2.3) the likelihood may be implicitly defined; nevertheless, Bayes’ theorem underlies those analyses as well.




Figure 2.1.: An illustration of the difference between marginalizing and conditioning for a multivariate Gaussian distribution.

A set of N samples² $\{\theta_i^{(k)}\}_k$ can be used to represent a high-dimensional posterior distribution. This is convenient, since marginal distributions are trivially obtained from these by only considering one of their entries, since $\{\theta_i^{(k)}\} \sim p(\theta_i|d)$.

We can easily compute a number of summary statistics from samples, such as the expectation value of any function $f(\theta)$ [178]:

$$\langle f \rangle_{\theta \sim p(\theta|d, M)} \approx \frac{1}{N} \sum_{k=1}^N f(\theta^{(k)}).$$

This general expression can be used for many different applications:

- the mean and variance of a parameter θ_i can be obtained by setting $f = \theta_i$ and $f = (\theta_i - \langle \theta_i \rangle)^2 N / (N - 1)$ respectively;
- the covariance matrix of the whole parameter space can be obtained by setting $f = (\theta_i - \langle \theta_i \rangle)(\theta_j - \langle \theta_j \rangle) N / (N - 1)$;
- a histogram for the parameter θ_i can be obtained with a set of functions which return the height of bin m : $f_m = [\theta_m^{\min} \leq \theta_i^{(k)} < \theta_m^{\max}]$,³.

Working with samples also has disadvantages. For example, conditional distributions, which we are sometimes interested in, are difficult to obtain. Furthermore, in certain applications an estimate of the underlying density function is necessary. Techniques to obtain it exist, including Kernel Density Estimation and normalizing flows (Section 2.3.1), but they are not closed-form computations and often require tuning.

2.1.1.1. Credible regions

A common method to quantify the uncertainty in the estimate of a parameter is to compute some kind of Bayesian Credible Interval with a given confidence X ; in the gravitational wave literature, $X = 0.9$ is generally the choice.

There are different ways to define such an interval, depending on the problem at hand. If a parameter is bounded, and its posterior exhibits significant mass near the boundary, a single-sided interval is often given:

$$“\theta_i \leq y \text{ at } X \text{ confidence}” \iff \int_{\theta_i^{\min}}^y p(\theta_i|d) d\theta_i = X$$

²We represent these samples with two indices: a lower one indicating the dimension in the parameter space, and an upper one in parenthesis indicating the number of the sample.

³This expression is using the Iverson bracket [106] which returns 1 if the statement inside is true and 0 otherwise.

$$“\theta_i \geq y \text{ at } X \text{ confidence}” \iff \int_y^{\theta_i^{\max}} p(\theta_i|d)d\theta_i = X .$$

If a parameter is not constrained near a boundary, a symmetric interval can be used:

$$“\theta_i \in [y_1, y_2] \text{ at } X \text{ confidence, symmetric}” \iff \int_{y_1}^{y_2} p(\theta_i|d)d\theta_i = X ,$$

with also

$$\int_{y_2}^{\theta_i^{\max}} p(\theta_i|d)d\theta_i = \int_{\theta_i^{\min}}^{y_1} p(\theta_i|d)d\theta_i = \frac{1 - X}{2} .$$

These intervals are commonly adopted due to the ease of their computation, which reduces to computing quantiles of the $p(\theta_i|d)$ distribution, either at X or $1 - X$ (for the one-sided case) or at $(1 - X)/2$ and $(1 + X)/2$ (for the symmetric case).

Another method to quantify the uncertainty on one or more parameters is the *Highest Probability Density* (HPD) region. This is a region defined as $R(X) = \{\theta | p(\theta|d) \geq K(X)\}$, where the value of K is chosen such that

$$\int_{R(X)} p(\theta|d)d\theta = X .$$

For a univariate and unimodal distribution, this is the shortest possible interval which contains a probability mass X , and it can be computed quickly by sorting the samples, computing all possible interval widths and choosing the minimum.

However, this technique will fail with no warning if the posterior distribution is not unimodal, since then the HPD may consist of disjoint intervals. A computationally expensive, but more accurate way to estimate this region is:

- discretizing the parameter space for the variables of interest into bins;
- estimating the probability density in each bin (this is often performed with a KDE, since counting the samples in each bin is affected by high variance);
- sorting the bins by contained probability mass (highest first);
- adding the probability mass in the first n bins, and increasing n until the total surpasses X .

The n bins obtained will represent the desired region.

2.1.1.2. Weighted samples

Some algorithms, such as Nested Sampling (Section 2.2) or importance-sampled Neural Posterior Estimation (Section 2.3.3.3), can return a set of *weighted* samples. This means that each sample θ_i is associated with a weight w_i . Weighted samples can also arise in the context of *reweighting*: slightly perturbing the assumptions underlying an analysis and thus obtaining a set of weights in the form

$$w_i = \frac{\mathcal{L}_1(d|\theta_i)\pi_1(\theta_i)}{\mathcal{L}_2(d|\theta_i)\pi_2(\theta_i)}. \quad (2.2)$$

Here, it is assumed that the samples θ_i were obtained through an analysis with model 2, so they are distributed according to the posterior distribution $p_2(\theta_i)$.

The weights w_i could in principle also be defined as the ratio of the two normalized posteriors, multiplying them by a factor Z_2/Z_1 . However, when reweighting we typically only have access to an estimate of the evidence Z_2 from the previous analysis, while Z_1 still to be determined. Defining the weights without this term, as in Equation 2.2, means that their average will provide this estimate (see also Section 2.3.3):

$$\frac{1}{N} \sum_{i=1}^N w_i \approx \frac{Z_1}{Z_2}.$$

Summary statistics can then be computed through **importance sampling**, as

$$\langle f \rangle_{\theta \sim p(\theta|d,M)} \approx \frac{\sum_{i=1}^N w_i f(\theta_i)}{\sum_{i=1}^N w_i}. \quad (2.3)$$

Based on this, all the quantities that could be obtained as described in Section 2.1.1 can be retrieved with weighted samples as well.

We can also extract equally-weighted samples from weighted ones by **rejection sampling** [see *e.g.* 10, for a review]:

1. Establish the highest weight, $w_{\max} = \max w_i$;
2. for each sample θ_i :
 1. compute $x = w_i/w_{\max} \in [0, 1]$;
 2. extract a uniformly random number $r \in [0, 1]$;
 3. if $r < x$, keep the sample θ_i , otherwise, remove it.

On average, this procedure will retain $n\langle w \rangle / \max w$ out of n samples; if many of the weights are lower than their maximum, it can become quite inefficient. In fact, under the realistic assumption of the weights being log-normally distributed, it can be shown that working with all weighted samples can be significantly more efficient than performing rejection sampling and then working with the resulting equally-weighted samples [10, fig. 3].

2.1.2. p - p tests

Mistrusting the results of our analysis pipelines is wise: we ought to put them through as many tests as possible in order to spot any issues. A common test for self-consistency is the p - p plot, which specifically regards the marginal distributions, $p(\theta_i|d)$ for a fixed i .

This test starts with a campaign of injections, for which the true value $\hat{\theta}_i$ is known: every time, we need to generate a realistic simulation of the data d and analyze it, feigning ignorance of the true value. Then, we may compute the percentile of the distribution at which the true value lies:

$$u = \int_{-\infty}^{\hat{\theta}_i} p(\theta_i|d) d\theta = \text{CDF}_{p(\theta_i|d)}(\hat{\theta}_i). \quad (2.4)$$

Here we assume for simplicity that the parameter can range over $\mathbb{R}$, but the same calculation can be performed starting integration from the lower bound. This is the definition of the cumulative distribution function, computed at $\hat{\theta}_i$. Repeating this procedure N_{inj} times, we will get a set of values u_j , where the index j denotes the number of the injection. If our analysis is self-consistent, and $p(\theta_i|d)$ is the probability density for the true value, then the u_j should be uniformly distributed in the interval $[0, 1]$.

There are a few ways to check whether this is the case, both visual and statistical. The simplest is probably to create a histogram of the u_j and visually check for uniformity. This is not a very powerful method, however, since it is hard to tell by eye which level of fluctuation is “acceptable”, and there is a dependence on the size of the bins used for the histogram. A better choice is to compute the empirical cumulative distribution function:

$$P_{\text{emp}}(u) = \frac{\text{number of } u_j \leq u}{N_{\text{inj}}}$$

This is a “stepwise”, locally constant function. For a perfectly uniform distribution, and an asymptotically high number of injections, it asymptotically tends toward $P(u) = u$. It is uniquely defined, with no parameters.

A first visual test is to plot it and check that it approximately follows the diagonal. A useful variant of this is to plot u on the horizontal axis and $P(u) - u$ on the vertical, which

is expected to be near zero; see e.g. the work of Chatziioannou et al. [44]. This allows for the resolution of small fluctuations from the diagonal.

In order to help in the visual check, we need something to compare against: the magnitude of the expected fluctuations of $P(u)$ under the null hypothesis, i.e. if the distribution is indeed uniform. If we fix a value of u , this can be reframed as counting, out of N_{inj} independent and identically distributed draws from the distribution $\mathcal{U}([0, 1])$, how many are smaller than u . The probability for each of them being smaller than u is u itself, so the probability mass function describing this scenario will be binomial, but scaled by N_{inj} :

$$p\left(P(u) = \frac{k}{N_{\text{inj}}}\right) = \binom{k}{N_{\text{inj}}} u^k (1-u)^{N_{\text{inj}}-k}$$

Therefore we can easily compute its mean and variance: $\langle P(u) \rangle = u$ and

$$\langle (P(u) - u)^2 \rangle = \frac{u(1-u)}{N_{\text{inj}}}.$$




Figure 2.2.: Visualizations for a p-p test: histogram, empirical cdf, tilted empirical cdf. We also visualize the two modes of the Kolmogorov-Smirnov test. This plot was inspired by the visualizations in Williams, Veitch, and Messenger [195].

Along with the visual tests, we can perform a Kolmogorov-Smirnov test, which can give us a p -value computed under the null hypothesis of the distribution being uniform.

In Figure 2.2 we show these for the one-sided Kolmogorov-Smirnov [108, 167] statistics D_+ and D_- , defined in terms of the empirical CDF as

$$D_{\pm} = \sup_u [\pm (P_{\text{emp}}(u) - u)] .$$

We will always have $D_{\pm} \geq 0$, since by definition $P_{\text{emp}}(0) = 0$, meaning that the set of values we are taking the sup of includes 0. One may also use the two-sided statistic

$$D = \sup_u |P_{\text{emp}}(u) - u|.$$

We show both one-sided tests as opposed to the two-sided one, since their values can be used together to distinguish different types of misspecification, as we illustrate in Section 2.1.3. The density function of the K-S statistics under the null hypothesis is shown in Figure 2.3. The densities for the $D_{\pm}$ functions are identical - the transformation $u \rightarrow 1 - u$ maps one to the other, and the null hypothesis is invariant under this reflection. The scenario $D \approx 0$ is unlikely, as it would require the draws to precisely align themselves in an equally spaced pattern. On the other hand, $D_{\pm}$ being close to zero is plausible, as it only requires the realization to mostly be above or below the uniform CDF.




Figure 2.3.: Distribution for the K-S statistics.

These tests are common in the literature: they were performed, for example, for the software BILBY [151]. More details on them can also be found in Romero-Shaw, Thrane, and Lasky [152]. Within the sampler nessai [195], a p - p plot test is applied to the insertion indices of new nested samples, *i.e.* the index of the new sample when ranking all previous live points by likelihood. As we shall discuss in Section 2.2, the new sample should be drawn uniformly in the likelihood-constrained prior mass. If this is the case, and the previous samples are also uniform in the variable X (which denotes prior mass), then this insertion index should be uniformly distributed between 0 and n_{live} , which can be investigated with a p - p plot. Note that in that case, overconstraining of the distribution being sampled will not correspond to overconstraining in terms of *insertion indices*, but instead be equivalent

to an underestimation bias, with more quantiles near the top-end of the distribution.⁴

2.1.3. p - p test failure modes

Let us now explore the way that p - p plots look when the test fails. We will assume the true parameter θ is always equal to zero, the noise is Gaussian, and our prior on the parameter θ is flat. This will allow us to explore how the p - p plot can help us diagnose errors in the estimated posterior distribution.

For all the cases considered, we simulate injections analytically as follows:

1. we draw a noise instance, which we assume to be normally distributed as $n \sim \mathcal{N}(n; \mu = 0, \sigma = 1)$, the observed value will be $n + \hat{\theta} = n$;
2. we define a posterior distribution $p(\theta|d)$ — the correct one would be $p(\theta|d) = \mathcal{N}(\theta; \mu = n, \sigma = 1)$, but we will also use other (incorrect) options;
3. we compute the quantile u based on $p(\theta|d)$ and $\hat{\theta} = 0$ as in equation Equation 2.4.

First, let us suppose we have a bias in our measurement, so that the posterior distribution is not centered at the measured value, $\hat{\theta} + n = n$, but instead at a slightly higher one: we will use a small deviation, 0.2. Our posterior distribution will then read

$$p(\theta|d) = \mathcal{N}(\theta; \mu = n + 0.2, \sigma = 1).$$

The result is in Figure 2.4.




Figure 2.4.: p - p plot in the scenario of an overestimation bias.

⁴Also note that the nessai diagnostic plots are set up to show a quantity defined with an opposite sign compared to this work, equivalent to the difference $u - P(u)$.

Figure 2.5 shows what happens if the bias is in the negative direction:

$$p(\theta|d) = \mathcal{N}(\theta; \mu = -0.2 + n, \sigma = 1) .$$




Figure 2.5.: p-p plot in the scenario of an underestimation bias.

Note that only one of the one-sided KS tests gives extreme results in either case.

Now let us consider a different scenario: when the mean of the posterior distribution is correctly estimated, but its variance is not. If the variance is too high, we are *underconstraining* the parameter, assigning a higher uncertainty to our measurement than is warranted: for example, if

$$p(\theta|d) = \mathcal{N}(\mu = n, \sigma = 1.2) .$$




Figure 2.6.: p-p plot in the scenario of an underconstraining bias.

Figure 2.6 shows the result. In the opposite scenario, we are *overconstraining* (Figure 2.7), assigning a lower uncertainty to our measurement than is warranted: for example,

$$p(\theta|d) = \mathcal{N}(\mu = n, \sigma = 0.8).$$




Figure 2.7.: p-p plot in the scenario of an overconstraining bias.

Finally, let us consider a distribution with the correct mean and variance, but the wrong shape. To this end we choose a generalized normal distribution, with a shape parameter $\beta = 2.5$. The $\beta = 2$ case recovers the standard normal distribution.




Figure 2.8.: Generalized normal and normal distributions with the same variance and mean.

We rescale the distribution so that it has unit variance, as shown in figure Figure 2.8. As this shift is harder to detect than the ones before, we need to raise the statistics.




Figure 2.9.: p-p plot in the scenario of the wrong distribution, despite correct mean and variance.

The distribution looks similar to the underconstrained case, as for the bulk of the probability mass (-2 to +2 in Figure 2.8) the difference between these distribution is analogous to underconstraining. However, close to the edges in Figure 2.9 we can see the trend reverse.

2.1.4. Model comparison

The evidence $\mathcal{Z}$ is helpful when comparing models. In a realistic scenario, our model M may not adequately capture the behavior of the system at hand, but in Equation 2.1, all the terms are conditional on it: an analysis performed under the assumption that the model is true cannot inform us on its validity, by construction. If more than one model is available, we can make some steps toward addressing this problem by comparing them; still, we should remember that our results are dependent on the set of models we used, and it is possible they may *all* miss some important aspect.

We can write an analogous expression to Equation 2.1 for both models M_1 and M_2 :

$$p(M_i|d)Z(d) = \mathcal{Z}(d|M_i)\pi(M_i);$$

dividing one by the other, we get the following:⁵

$$\frac{p(M_1|d)}{p(M_2|d)} = \frac{\mathcal{Z}(d|M_1) \pi(M_1)}{\mathcal{Z}(d|M_2) \pi(M_2)}. \quad (2.5)$$

The left-hand term in Equation 2.5 is dubbed *posterior odds*: it describes the degree to which we should prefer model 1 over model 2 after seeing the data d . On the right hand we have a product of two terms: the *Bayes Factor* (BF)

$$\text{BF}_2^1 = \frac{\mathcal{Z}(d|M_1)}{\mathcal{Z}(d|M_2)} \quad (2.6)$$

describes the degree to which the data prefer model 1 over model 2. The remaining term, $\pi(M_1)/\pi(M_2)$, is called *prior odds*, and it describes the degree to which we preferred model 1 over model 2 before seeing the data d . This may be informed, for example, by previous observations, theoretical considerations, or simply be set to 1 if our belief is that both models are equally likely.

i Note 1: On the usage of odds

Equation 2.6 can be applied several times in a row, as more observations become available; the Bayes Factors from each observation multiply.

This has an interesting consequence if the two models at hand are complementary, i.e. one being correct excludes the other, and one or the other must be correct, so that their probabilities satisfy $p(M_1) + p(M_2) = 1$. If this is the case, we may simplify the notation by simply using $P(M_1)$ and $1 - p(M_1) = p(M_2)$. Then, taking the logarithm of Equation 2.6, we can express it as

$$\text{logit}(p(M_1|d)) = \log \text{BF}_2^1 + \text{logit}(p(M_1)),$$

where we defined the logit function, also shown in figure Figure 2.10:

$$\text{logit}(p) = \log \frac{p}{1-p}.$$

⁵While both are probability distributions over the possible data realizations, we denote with $\mathcal{Z}$ the conditional evidence under the assumption of model M_i , while Z is the marginal probability of the data without the assumption of this model. One might interpret it as the likelihood of the data under the common assumptions shared by all models M_i , e.g. regarding the noise-generating process, which all our results are implicitly conditioned over. This is anyway immaterial, since the $Z(d)$ term cancels in the equation.




Figure 2.10.: The logit function.

The logit of a probability, $\text{logit}(p) \in \mathbb{R}$, also called **log-odds**, can be argued to be a more natural way to express it than $p \in (0, 1)$. Probabilities very close to 0 or 1 are extreme statements: in this parameterization, they are accordingly very far from the origin, and absolutes such as $p = 0, 1$ are divergent towards $\pm\infty$.

Based on this formalism, we may be tempted to multiply Bayes Factors from different observations of the same phenomenon. While this is correct in an idealized scenario, it can lead to dramatically incorrect conclusions when applied in practice, due to its dependence on the prior used in each analysis; Isi, Farr, and Chatziioannou [99] have discussed this in the context of tests of General Relativity (GR).

If two models are nested, *i.e.* one is a special case of the other, then we can express the Bayes Factor between them in a simple way [66, 39]:

$$\text{BF}_1^2 = \frac{p_2(\theta = \theta_0|d)}{\pi_2(\theta = \theta_0)}. \quad (2.7)$$

Here, model 1 is the special case of model 2 in which parameter θ equals θ_0 , p_2 is the marginal posterior distribution for the parameter θ under model 2, while π_2 is the marginal prior distribution for the parameter θ under model 2.

The value of π_2 computed at any given point is sensitive to how wide the prior is, since it must be a normalized probability distribution. For example, the density of a uniform prior over an interval of width A is precisely $1/A$.

This formulation illustrates the general fact that Bayes Factors are heavily dependent on the width of the prior distribution. This is a well-known feature of Bayesian inference, known as the Occam Factor or Penalty. Among two models which fit the data equally well (*i.e.* the maximum values of their likelihoods are equal), the preferred one is the one for which the likelihood has high values for a larger fraction of the prior space. Qualitatively, this penalizes flexible models, which are able to fit many different scenarios beside the observed one, in favor of ones which only fit the observed scenario and few others.

A negative consequence of this arises if the prior distributions cannot be selected in a principled way, though: if they are set too wide, the data may seem to favor the baseline model (*i.e.* model 1 in equation Equation 2.7) even in the presence of a deviation (model 2 in the same equation with $\theta \neq \theta_0$). By the way Bayes Factors are constructed, if the BF for any single observation favors the baseline model, combining many of them from different observations only accumulates evidence against a deviation, even if a small deviation is present and could be detected by a simultaneous analysis of all events.

2.2. Nested sampling

Nested sampling, introduced by Skilling in the early 2000s (Skilling [165], Skilling [166]), is an algorithm designed to compute the *evidence* $\mathcal{Z}$ as well as the posterior distribution on the parameters.

The evidence is computed through the integral

$$\mathcal{Z}(d|M) = \int \mathcal{L}(d|\theta, M) \pi(\theta|M) d^n \theta.$$

This integral is in n dimensions (where n is the dimensionality of the parameter space) and therefore difficult to work with. The idea in nested sampling is to rewrite it in terms of the auxiliary variable X , defined as the prior volume contained within a likelihood constraint $\mathcal{L} \geq L$:

$$X(L) = \int_{\mathcal{L}(d|\theta, M) \geq L} \pi(\theta|M) d^n \theta.$$

Since the prior is a normalized probability distribution and the likelihood ranges from 0 to $L_{\max}$, this variable will range from $X(0) = 1$ to $X(L_{\max}) = 0$, and it will always be

decreasing.⁶ We can define its inverse, $L(X)$, and rewrite the likelihood integral through integration by parts: the evidence is the expectation value of the likelihood, therefore, starting from the definition of a Lebesgue integral:

$$\mathcal{Z}(d|M) = \int_0^{L_{\max}} X(L) dL = XL|_{L=0}^{L=L_{\max}} - \int_{X(L=0)}^{X(L=L_{\max})} L(X) dX = \int_0^1 L(X) dX. \quad (2.8)$$

For a detailed discussion of the mathematical details, also see Ashton et al. [9], Sivia and Skilling [164].

2.2.1. Prior compression

Note that, by definition, X is directly related to probability mass: the prior probability of an interval $[X_0, X_1]$, i.e. the prior mass which has likelihood values $L(X_1) < \mathcal{L} < L(X_0)$, is

$$p([X_0, X_1]) = \int_{L(X_0) < \mathcal{L}(d|\theta, M) < L(X_1)} \pi(\theta|M) d^n \theta = X_1 - X_0,$$

where we computed the integral on a likelihood “ring” by subtracting the inside volume from the outside one. This is useful to us, since it means that if we distribute points according to the prior their X values will also be uniformly distributed in $[0, 1]$. More formally, the probability we are assigning to X is the *push-forward* probability measure associated to the mapping from θ to X defined by $X(\theta) = X(L(\theta))$. For a discussion on this, see [9, box 2].

If k points are uniformly distributed in an interval $[0, X^*]$, then the largest of them will have a X coordinate distributed as $X_{\max}/X^* \equiv t \sim \text{beta}(k, 1) = kt^{k-1}$, where we defined the *compression ratio* t : when we discard the lowest-likelihood point with $X_{\max}$, we will have contracted the prior volume by a factor t .

i Note 2: On the beta distribution

Here is a simple argument for why t , the largest out of k uniform random variates in $[0, 1]$, is distributed according to $t \sim kt^{k-1}$. The statement can equivalently be framed in terms of the cumulative distribution: $p(t \leq T) = T^k$.

We can think of our k random variates as a point in the k -dimensional unit box $[0, 1]^k$,

⁶It will not necessarily be *strictly* decreasing, but that is not a conceptual issue. One can introduce jitter in the likelihood, varying it at each point by an inconsequential amount, in order to ensure that the decreasing condition is verified precisely.

and since they are uniformly distributed there will be a one-to-one correspondence between volume within the box and probability mass. Then, the geometric meaning of the cumulative distribution is the volume of the region corresponding to the event $t \leq T$, i.e. such that the maximum of the coordinates of our point is T . The answer is the box $[0, T]^k$, whose volume is T^k , thus proving our claim.

In general, the n -th smallest out of k uniform random variates in $[0, 1]$ is distributed according to a beta distribution, specifically $\text{beta}(k, n - k + 1) \propto x^{n-1}(1-x)^{k-n}$; here we are looking at the k -th smallest (i.e. the largest).

In terms of $\log t$, the distribution reads

$$\frac{dp}{d \log t} = t \frac{dp}{dt} = kt^k = ke^{k \log t}$$

This distribution is shown in Figure 2.11.




Figure 2.11.: Volume compression

Its mean and standard deviation are:

$$\begin{aligned}
\langle \log t \rangle &= \int_{-\infty}^0 \log t \frac{dp}{d \log t} d \log t \\
&= \int_{-\infty}^0 \frac{x}{k} k e^x \frac{dx}{k} \\
&= \frac{1}{k} \int_{-\infty}^0 x e^x dx \\
&= -\frac{1}{k}
\end{aligned}$$

where $x = k \log t$, and

$$\begin{aligned}
\langle (\log t)^2 \rangle - \frac{1}{k^2} &= \int_{-\infty}^0 (\log t)^2 \frac{dp}{d \log t} d \log t - \frac{1}{k^2} \\
&= \int_{-\infty}^0 \frac{x^2}{k^2} k e^x \frac{dx}{k} - \frac{1}{k^2} \\
&= \frac{1}{k^2} \int_{-\infty}^0 x^2 e^x dx - \frac{1}{k^2} \\
&= \frac{2-1}{k^2} = \frac{1}{k^2}.
\end{aligned}$$

Compactly, we can then write $\log t = (-1 \pm 1)/k$.

2.2.2. Computing evidence and posterior

If we sample k points from the prior (with total volume X_0) and remove the one with the worst likelihood, the remaining prior volume will be $X_1 = t_1 X_0$, where $\log t_1$ has an expectation value of $-1/k$. This will be true at every iteration, so [166, section 5]

$$\log X_i = \sum_{j \leq i} \log t_j \approx -\frac{i}{k} \pm \frac{\sqrt{i}}{k}.$$

The prior is compressed exponentially, which is desirable since the typical set for the posterior is often several orders of magnitude smaller than the prior. For a visualization of this compression see Figure 2.12.

This allows us to construct a procedure to evaluate the integral in Equation 2.8, as long as we are able to do the following:




Figure 2.12.: Volume compression by iteration. The number of points is denoted here as $n_{\text{live}} = k$. Note how horizontal slices of this ensemble are thinner when we have more live points: increasing them will decrease the uncertainty at a fixed prior compression. We illustrate this by drawing a line corresponding to the one-sigma range of the distribution of prior compressions obtained after $4k$ iterations. After reaching a logarithmic volume compression of $-\log X = 4$ the runs are truncated, and the remaining points are shown: these are distributed uniformly and not geometrically in X , therefore they taper off in this logarithmic plot.

1. sample k points from the prior;
2. find the lowest-likelihood point θ_1 and discard it: now the prior volume is approximately $X_1 = e^{-1/k}$, let us denote its likelihood as L_1 ;
3. sample a new point θ_2 from the prior, constrained to $L_2 = L(\theta_2) > L_1$;
4. repeat from point 2.

The points which still lie above the likelihood threshold at any given iteration are called *live* points. Since every time one is discarded we promptly replace it, their total number k (also often written as n_{live}) is conserved. We are constructing a sequence of discarded (“dead”) points θ_i with corresponding likelihoods L_i and approximate associated interior prior volumes $X_i \approx e^{-i/k}$.

At this stage, we still need to discuss how step 3 is performed, as well as when the iteration should stop — for now, suppose that happens at some stage n_{iter} .

For the evidence, we can simply use a trapezoidal integration rule:

$$\mathcal{Z}(d|M) \approx \sum_{i=1}^{n_{\text{iter}}} w_i L_i = \sum_{i=1}^{n_{\text{iter}}} \frac{X_{i+1} - X_{i-1}}{2} L_i$$

It can be shown that this sum converges to the correct value in the $n_{\text{iter}} \rightarrow \infty$ and $n_{\text{live}} \rightarrow \infty$ limit.

These weights can also be used to estimate the posterior as follows: take the points θ_i and assign to each of them a weight

$$p_i = \frac{w_i L_i}{Z}. \quad (2.9)$$

These weighted samples (see Section 2.1.1.2) represent the posterior. It is readily seen that these weights add to 1, and heuristically they are computed as one might expect, by multiplying the prior mass and likelihood value for each of the regions we have subdivided the prior volume into.

2.2.3. Sampling options

Nested sampling is, in a way, a meta-algorithm: the sub-algorithm used to sample a new point from the prior subject to the constraint that its likelihood be higher than the last rejected point is a crucial determinant of the computational complexity and behavior of the implementation. This problem is called Likelihood-Restricted Prior Sampling.

A review of several possible options can be found in [35, section 5] or [9, table 2]. Roughly speaking, they can be classified into:

- MCMC-based methods;
- region-based methods;
- machine learning-based methods.
- hybrid methods;

While the current set of live points is often used as a guide, the statistical properties of nested sampling hold as long as the new point is uncorrelated from the previous ones.

2.2.3.1. The prior transform

The way most nested sampling software implements likelihood-restricted prior sampling starts from the assumption that the prior is uniform in the box $[0, 1]^d$.

We can, however, use arbitrary non-uniform priors: for any given prior $\pi(\theta)$ we want to use, we need to find a mapping $g: [0, 1]^d \rightarrow \mathbb{R}^d$ between the unit d -dimensional box and our physical prior space, such that if $u \sim \mathcal{U}([0, 1]^d)$, then $g(u) \sim \pi$, i.e. mapping uniformly-distributed random numbers onto values distributed according to our prior.

If our prior is *separable*, i.e. $\pi(\theta) = \prod_{i=1}^d \pi_i(\theta_i)$, then the problem is solvable analytically: the prior transform map will be such that

$$g(u) = [\text{CDF}_{\pi_i}^{-1}(u_i)]_i$$

where “CDF” denotes the cumulative distribution of π_i :

$$\text{CDF}_{\pi_i}(\theta_i) = \int_{-\infty}^{\theta_i} \pi_i(y) dy \in [0, 1]$$

and CDF^{-1} is its inverse. The cumulative distribution can also be computed from samples, so in the case of separate priors we need not have an analytical expression for them.

In the non-separable case, while it may be mathematically possible to define the transform g corresponding to the distribution π , there is no general procedure to obtain it; equivalently, there is no algorithm to “analytically” sample from an arbitrary distribution π .

It is still possible to work with non-separable priors: for example, one could model it with a normalizing flow (Section 2.3.1), or if its density function is known one could perform the analysis with an approximate analytic prior and then reweight the samples as described in Section 2.1.1.2. These approaches are, however, subject to potential pitfalls (failed convergence, low sample efficiency, dependence on hyperparameters), and their performance needs to be studied on a case-by-case basis.

2.2.4. Model comparison with nested sampling

In order to illustrate the process of model comparison, let us define a simple problem: suppose we have a set of data points (x_i, y_i) with constant, known and Gaussian uncertainties σ_y on the y values.

We will compare two simple *nested* models, namely a parabolic one with $y = f(x; a, b, c) = ax^2 + bx + c$ and a linear one with $y = f(x; b, c) = bx + c$. Model comparison can also be performed for non-nested models, but the nested case has some nice properties we will showcase later.

Figure 2.13 shows the set of data we will use, which comes from the parabolic model.




Figure 2.13.: Toy model, with data coming from a parabolic function and Gaussian errors.

We can then analyze this data, using a Gaussian likelihood and setting (somewhat arbitrarily) uniform priors for the parameters a , b and c in the range $[-10, 10]$:

- $\pi(a, b, c) = 1/20^3$ in that range and 0 otherwise for the parabola;
- $\pi(b, c) = 1/20^2$ in that range and 0 otherwise for the parabola;

The log-likelihood will be

$$\log \mathcal{L}(\text{data}|\theta, \text{model}) = -\frac{1}{2} \sum_{i=1}^n \frac{(y_i - f(x_i; \theta))^2}{\sigma_y^2} - \frac{n}{2} \log(2\pi\sigma_y^2).$$

We perform Bayesian inference for the two models with Nested Sampling (Section 2.2). It gives us estimates for the posterior distributions of the parameters as well as the evidences. This allows us to compute a Bayes Factor:

$$\log \text{BF}_{\text{line}}^{\text{parabola}} = \log Z_{\text{parabola}} - \log Z_{\text{line}} ,$$

which appears in the equation for the posterior odds for one model against the other:

$$\underbrace{\frac{p(M_{\text{parabola}}|d)}{P(M_{\text{line}}|d)}}_{\text{posterior odds}} = \text{BF}_{\text{line}}^{\text{parabola}} \underbrace{\frac{p(M_{\text{parabola}})}{p(M_{\text{line}})}}_{\text{prior odds}} .$$

In this case, the curvature of the data is strong enough for the Bayes Factor to favor the parabolic model (although not by much): the log Bayes factor is 2.12 ± 0.11 nats in favor of the parabolic model.

Since the models we are considering were nested, we have a comparison point for our Bayes Factor: the Savage-Dickey density ratio; see Equation 2.7.

The linear model is indeed a special case of the parabolic one (with $a = 0$); the priors for the other parameters are exactly the same in both cases. Then:

$$\text{BF}_{\text{line}}^{\text{parabola}} = \frac{p(a = 0|d, M_{\text{parabola}})}{\pi(a = 0|M_{\text{parabola}})} .$$

One might then think that, at least in this type of scenario, going through the evidence computation with a powerful tool such as Nested Sampling was unnecessary: the posterior distribution can be computed for cheaper, after all. However, as we show in Figure 2.14, computing the value for the posterior density distribution at the point $a = 0$ is not easy.

We can approximate the posterior density based on the samples with tools like a histogram or a kernel density estimate, but any such method will suffer from low statistics if we need to explore the edges of the distribution, as we see here: we would need an enormous amount of posterior samples to reach the same accuracy on the ratio of the distributions we get with Nested Sampling.

2.2.5. Commonly used plots

The quantities we discussed can be tracked through a Nested Sampling run as a useful diagnostic. Here we show some commonly used plots, based on a two-dimensional example where the likelihood is Gaussian in the parameters. The assumptions are as follows:




Figure 2.14.: Consistency of the Savage-Dickey density ratio and Nested Sampling estimates for the Bayes Factor. The blue curve is the prior on a under the parabolic model, the dark and light red curves are respectively a KDE reconstruction and histogram of the posterior distribution on a under the parabolic model. The dashed vertical line is the true value of a . Finally, the black dot represents the equivalent posterior-to-prior ratio to the Bayes Factor obtained through nested sampling. The value is broadly consistent with what one can obtain by approximating the posterior, but it is much more precise and its error is better quantified.

- our prior is a 2-dimensional uniform distribution $\pi \sim \mathcal{U}([0, 1]^2)$;
- our likelihood is a 2-dimensional normal distribution, with mean $\mu = [0.5, 0.5]$ and correlated covariance matrix

$$C = 3 \times 10^{-2} \begin{pmatrix} 1 & 0.95 \\ 0.95 & 1 \end{pmatrix}$$

We use 400 live points, and stop iterating when the contribution to the integral from a single iteration reaches $\Delta \log Z_i < 0.01$. We use the dynesty sampler [168].

2.2.5.1. Run plot

The first diagnostic tool is the *run plot* (Figure 2.15), where we see the accumulation of the evidence integral as a function of the prior compression $\log X$.

- the number of live points, which is constant up to the point where the nested sampling run is stopped: then, the remaining points are reused;
- the likelihood L_i , which is strictly increasing as expected;
- the posterior mass $w_i L_i / Z$, which increases and then decreases: first, the likelihood's increase dominates, but at some point its values start tapering, and the decrease in prior volume w_i dominates;
- the accumulation of evidence, with a corresponding error estimated as discussed in Section 2.2.6.

2.2.5.2. Trace plot

In Figure 2.16 we show a *trace plot*, which relates the evolution in prior compression to the actual values of the parameters. For each parameter, we get a scatter plot of its value as a function of the prior compression, and we can see its range shrink as we reach higher and higher values of the likelihood.

More specifically, this plot is showing the locations of the *dead points*, i.e. the lowest-likelihood points that get discarded at each iteration. When one of them gets replaced, it will be assigned a new likelihood in the range $[L_i, L_{\max}]$, and after a certain number of iterations it will be replaced again. This is what the red lines are showing; they should be moving “randomly” across the parameter space, if we were to see them remain in a given region in the parameter space we would have an indication that our replacement algorithm is not performing properly.

The points are colored by their importance weights, $w_i L_i / Z$, and to the right we see a density estimate plot for the posterior distribution of each parameter together with a 95% confidence interval.




Figure 2.15.: A run plot, illustrating the progress of a nested sampling run.




Figure 2.16.: A trace plot, illustrating the parameter constraints for a nested sampling run.

2.2.5.3. Corner plot

The *corner plot* (Figure 2.17) is a great tool to extract physical understanding about our model. It shows the marginal density plot for all our parameters, as well as for each pair of parameters.

2.2.5.4. Posterior predictive density plot

For this example the scenario introduced at the start of Section 2.2.5 is not adequate, so we will use the problem from Section 2.1.4.

A convenient visualization to gain insight into the question whether our model can reproduce the data is the *posterior predictive distribution*: the distribution expected of new data, conditional on our model and the data we observed.




Figure 2.17.: A corner plot, showing correlations in the posterior distribution.

$$\begin{aligned}
 p(y_{\text{pred}}|y_{\text{obs}}, M) &= \int p(y_{\text{pred}}, \theta|y_{\text{obs}}, M) d\theta \\
 &= \int p(y_{\text{pred}}|\theta, y_{\text{obs}}, M) p(\theta|y_{\text{obs}}, M) d\theta \\
 &= \int \underbrace{p(y_{\text{pred}}|\theta, M)}_{\text{model}} \underbrace{p(\theta|y_{\text{obs}}, M)}_{\text{posterior}} d\theta
 \end{aligned}$$

We can show this in different ways; an easy one to implement is by drawing samples from the posterior and plotting the model realization corresponding to each of them. The ensemble of curves will follow the posterior predictive distribution, as shown in Figure 2.18.

For an in-depth discussion of posterior predictive checks in the context of gravitational wave data analysis, see [152].

2.2.6. Information gain

The typical set of the posterior is often “much smaller” than the prior. A useful way to quantify this is the relative entropy, or Kullback-Leibler divergence between posterior and prior, often also called *information gain* or H :




Figure 2.18.: A posterior predictive density plot, sketching the posterior densities by plotting 100 curves sampled from the posterior.

$$H = \text{KL}(p \parallel \pi) = \int p(\theta) \log \left(\frac{p(\theta)}{\pi(\theta)} \right) d\theta = \int p(X) \log p(X) dX.$$

The base of the logarithms defines the unit used for the information gain:

- if they are natural logarithms (base e) the values are in *nats*;
- if they are base 2 logarithms, the values are in *Shannons* or *bits*;
- if they are base 10 logarithms, the values are in *Hartleys* or *dex*.

One way to interpret this is related to the signal-to-noise ratio [164, eq. 9.15]:

$$H \sim \# \text{ active components in the data} \times \log(\text{signal-to-noise ratio})$$

It can also be seen as a “volumetric compression” from the prior to the posterior. If the distributions are uniform, this intuition is exact [139]: suppose we have $\pi(\theta) \equiv [\theta \in V_\pi]/V_\pi$ and $p(\theta) \equiv [\theta \in V_p]/V_p$, where the two volumes are such that $V_p \subset V_\pi$.⁷ Then,

$$\text{KL}(p \parallel \pi) = \int_{V_\pi} \frac{[\theta \in V_p]}{V_p} \log \frac{V_\pi}{V_p} d\theta = \log \frac{V_\pi}{V_p},$$

⁷The square brackets are the Iverson bracket, as defined by Knuth [106].

The integral can be restricted to the V_p volume, since the integrand is zero outside it. Another useful special case to develop an intuition is the one in which our prior is a Gaussian distribution with standard deviation σ , and our posterior is centered on the same value but its standard deviation is $\sigma' = \sigma/k$. Then, the information gain is [34]:

$$\text{KL}(\mathcal{N}(\mu, \sigma') \parallel \mathcal{N}(\mu, \sigma)) = \log k + \frac{1}{2} \left(\frac{1}{k^2} - 1 \right).$$

We can compute the theoretical value for the entropy in the example above, since the computation is analytic in the Gaussian case: it is known that the differential entropy of a d -dimensional multivariate Gaussian random variable $f \sim \mathcal{N}(\mu, C)$ is [55, theorem 8.4.1]:

$$-\int_{\mathbb{R}^d} f(\theta) \log f(\theta) \, d^d\theta = \frac{d}{2} \log(2\pi e) + \log \det C.$$

The integral we need to compute H is very similar. The sign is opposite, and instead of being over $\mathbb{R}^2$ it is only over $[0, 1]^2$. The region outside of this box is, however, more than 10σ from the mean, therefore its contribution is negligible.

Comparing the analytical result with an estimate computed during the run by the sampler, we see that they are quite similar: the value of H obtained during sampling is 5.45, while the analytical estimate is 5.34.

2.2.6.1. Algorithm termination

The nested sampling algorithm we described above could continue forever, with ever-smaller compression $\log X$. However, if the likelihood function is bounded, after a certain number of iterations the prior volume will be concentrated around its maximum value, with all the live points having similar likelihood values, as Figure 2.19 shows.

Thus, we can find a termination condition by diagnosing whether this is happening: for example, we can compute the contribution to the evidence integral from the current live points, and stop iterating when that becomes smaller than a given threshold — this is the default in dynesty [168].

Since this is a stochastic algorithm, there is always a chance we are making a mistake when terminating: for example, there could be a region with very low volume but high likelihood that none of the live points have happened upon, but which could give a significant contribution to the evidence. There is no general solution to this issue, besides increasing the number of live points in order to improve the chances that some of them will happen upon the high-likelihood region.




Figure 2.19.: An illustration of the conditions for termination in nested sampling.

2.2.7. Acceleration

The information gain provides an estimate of the expected volumetric compression: we expect to be sampling from the bulk of the posterior mass when $\log X \sim -H$, which will take approximately $n_{\text{live}}H \pm \sqrt{n_{\text{live}}H}$ iterations due to Poisson variability.

If $n_{\text{live}}H$ is large, we expect the dominant source of uncertainty to be the one on the compression reached at that stage — it will be much larger than the uncertainty on the compression in the $\mathcal{O}(\sqrt{n_{\text{live}}H})$ iterations required to traverse the posterior mass. This leads us to the estimate

$$\text{std}(\log Z) \approx \sqrt{\frac{H}{n_{\text{live}}}}.$$

So, for a fixed variance on the evidence, we need to have $n_{\text{live}} \propto H$; therefore, the number of iterations is approximately $n_{\text{iter}} \propto n_{\text{live}}H \propto H^2$.

The time complexity for nested sampling can then be estimated as

$$T \approx n_{\text{live}} \times \langle t_{\text{like}} \rangle \times \langle n_{\text{replace}} \rangle \times H$$

where

- n_{live} is the number of live points used;

- $\langle t_{\text{like}} \rangle$ is the average likelihood evaluation time;
- $\langle n_{\text{replace}} \rangle$ is the average number of likelihood evaluations required to find a replacement for each dead point;
- H is the information gain between posterior and prior.

This equation assumes that likelihood evaluation is the primary computational requirement; this may not always be the case, especially for sophisticated proposal methods such as machine-learning based ones [195].

Based on this equation, we can think of several ways to accelerate sampling; Roulet and Venumadhav [154] have reviewed many of them in the context of gravitational wave parameter estimation.

The minimum number of live points required for unbiased inference is generally fixed by the problem, though as we discuss in Section 2.2.7.1 it may be varied during the run to gain in performance.

A large gain can be often obtained by improving the replacement efficiency of the method used — see also Section 2.2.3. Accelerating the evaluation of the likelihood, $\langle t_{\text{like}} \rangle$, will also result in a direct reduction of computational time. In Section 3.4 we discuss some approaches in this direction.

2.2.7.1. Varying the number of live points

The original nested sampling algorithm evolves a fixed number of live points throughout the run. It has since been proposed that it can be more efficient to vary it (Higson et al. [92], Speagle [168]), depending on the quantity we are more interested in evaluating.

A nested sampling run gives us estimates for both the evidence and the posterior. If what we are most interested in is precisely computing the evidence, we should allocate more points in the early stages of compression, even though the posterior mass is very low there: this will give us better statistics on the active volume when we reach the posterior bulk.

If we are more interested in the posterior, on the other hand, we can do the initial compression with fewer points and use more once we reach the posterior bulk.

It turns out that, while both these objectives can be optimized independently, the standard choice of using a constant number of live points is not on the Pareto front⁸ for the uncer-

⁸The Pareto front (introduced by Pareto [137] with the Italian name *serie di indifferenza*, “indifference series”) is a concept from multi-objective optimization. Defining multiple cost functions across a parameter space X establishes a partial order relation between points in X : if a point $x \in X$ has a lower cost than $y \in X$ for *all* cost functions, then we say that it strictly dominates y , or $x \succ y$. This is a partial ordering, since there are pairs (x, y) for which we have neither $x \succ y$ nor $y \succ x$. The *Pareto front* is then the set of points which are not dominated by any others.

tainties on evidence and posterior: a balanced dynamical allocation of points can improve both compared to the constant case [92].

2.3. Simulation-based inference

Simulation-based inference (SBI) encompasses a broad range of approaches to inference which do not require an explicit, analytic likelihood. While they still strive to produce results according to Bayes' theorem (Equation 2.1), they learn the likelihood implicitly, through the process of training a machine learning model. All SBI approaches require the presence of a *simulator*, i.e. a way to simulate data d conditional on a set of parameters θ , distributed according to the likelihood: $d \sim p(d|\theta)$. This need not be based on an analytic likelihood: for instance, in the gravitational wave case, one may inject signals into realistic detector noise [64, 71].

Several SBI methods exist, the main ones being [112]:

- **Neural Posterior Estimation** (NPE), the focus of Section 2.3.2, in which a density estimator $q(\theta|d)$ is trained to directly model the posterior density $p(\theta|d)$.
- **Neural Ratio Estimation**, in which a classifier is trained to distinguish between the joint distribution of parameters and data, $p(\theta, d)$ and the product of their marginals, $p(\theta)p(d)$. The optimal Bayesian classifier operates based on the decision function $f(\theta, d) = p(\theta, d)/(p(\theta, d) + p(\theta)p(d))$, from which the likelihood-to-evidence ratio $r(\theta, d) = p(d|\theta)/p(d) = f/(1 - f)$ can be easily recovered. This approach requires a stochastic sampling step.
- **Neural Likelihood Estimation**, in which a density estimator $q(d|\theta)$ directly learns the likelihood function. This approach also requires a stochastic sampling step.

Other approaches include Flow Matching Posterior Estimation and Consistency Model Posterior Estimation. The appeal of NPE specifically is in its promise to remove the stochastic sampling step, *i.e.* the requirement to use MCMC, nested sampling or other methods, thus providing an extreme gain in speed.

2.3.1. Normalizing flows

Modelling generic probability distributions is a task with wide applicability in gravitational wave science and beyond. A normalizing flow [136] is a series of parameterized transformations, whose joint action is denoted as T , which map a base distribution onto a generic other distribution. The base distribution should be easy to sample from: it is often chosen to be the standard n -dimensional Gaussian, with zero mean and the identity matrix as covariance. The transformations T are chosen to be invertible, and both T and T^{-1} to be

differentiable. Let us denote the variable in the “simple” space as u , and its transformed value in the parameter space as $\theta = T(u)$. If the density function for the simple distribution is labelled $p_u(u)$, its transformed version will be

$$q_\theta(\theta) = p_u(T^{-1}(x)) |\det J_{T^{-1}}(x)|.$$

Transformations in this class can be composed, and the determinant of the composition can be readily computed as the product of the individual determinants. This fact enables the flow-based architecture, in which several simple transformations are composed to achieve significant distortions of the base space. In fact, it can be shown under some weak assumptions that any distribution can be modeled in this manner.

If we have samples from a target distribution $p(\theta)$, we seek to train a normalizing flow $q_\theta(\theta; \lambda)$, where λ are the parameters defining the transform T . We can do so with the following loss function [136, equation 13]:

$$\begin{aligned} \text{loss}(\lambda) &= D_{\text{KL}}(p(\theta) || q_\theta(\theta; \lambda)) \\ &= \mathbb{E}_{\theta \sim p(\theta)} [-\log p_\theta(\theta; \lambda)] + \text{const.} \end{aligned} \tag{2.10}$$

A significant advantage of the normalizing flow architecture is that it enables evaluation of the density function $q_\theta(\theta)$ as well as the possibility to generate samples from it.

2.3.2. Neural posterior estimation

The objective of neural posterior estimation is to be able to directly model the posterior distribution based on an observation d . This is achieved through a network architecture dubbed a *neural density estimator* $q(\theta|d)$, which defines a normalized density function on the parameters θ conditioned on the data d . This is often a normalizing flow: see Section 2.3.1.

The loss function for this, in the simplest case, is [135, 61]:

$$\text{loss}_{\text{NPE}} = \mathbb{E}_{\theta \sim \pi(\theta)} \mathbb{E}_{d \sim p(d|\theta)} [-\log q(\theta|d)]. \tag{2.11}$$

This is equivalent, up to an additive constant, to the expectation value of the KL divergence between the posterior $p(\theta|d)$ and the estimator, $\mathbb{E}_{d \sim \mathcal{Z}(d)} D_{\text{KL}}(p(\theta|d) || q(\theta|d))$: in order to get Equation 2.11 from this expression, we need to apply Bayes’ theorem (Equation 2.1).

The expectation values can be numerically approximated by repeatedly sampling the parameters θ from their prior distribution $\pi(\theta)$, and then computing a set of data d by fixing the parameters and simulating a noise realization through the likelihood $p(d|\theta)$.

In practice, directly employing the loss in Equation 2.11 is impractical. Some parameters, such as the sky position of the signal, significantly alter the observed waveform in each detector by shifting it in time; This dependence, however, is known and can be amortized during training [61], leading to significantly improved performance. Additionally, the detector’s power spectral density S_n is varied during training, so that conditioning can be performed at inference time [194]. With this approach, the actual loss function for the state-of-the-art NPE model DINGO (“Deep Inference for Gravitational-wave Observations”) [64, equation 3] is

$$\text{loss}_{\text{GNPE}} = \mathbb{E}_{\theta \sim \pi(\theta)} \mathbb{E}_{S_n \sim p(S_n)} \mathbb{E}_{d \sim p(d|\theta)} \mathbb{E}_{\delta t_I \sim \kappa(\delta t_I)} [-\log q(\theta | d_{-t_I(\theta) - \delta t_I}, S_n, t_I(\theta) + \delta t_I)]. \quad (2.12)$$

Here, $t_I(\theta)$ is the arrival time of the gravitational wave signal at a given detector, and κ is a kernel used to perturb it. The subscript notation $d_{-t_I(\theta) - \delta t_I}$ denotes time-shifted data. This procedure makes it so that, in all the training data available to the model, waveforms are all approximately time-aligned.

Recently, a new transformer-based architecture for DINGO has been proposed [107] as an alternative to the GNPE approach described above. By tokenizing the data across detectors and frequency grid sections, a single network can be used for different detector configurations (one, two or three detectors), as well as for different frequency intervals.

2.3.3. Importance sampling

Taken alone, neural posterior estimation often yields fairly accurate estimates of the posterior distribution, but it does not do so in *every* case, and we have no guarantee that it did so for any *specific* signal of interest.

This is the motivation behind the introduction of neural importance sampling [62], which allows us to re-weight the posterior samples, estimate the evidence and correct the posterior distribution.

2.3.3.1. Importance sampling in general

Importance sampling is a common integration technique, based on the fundamental approximate equality

$$\int f(x)g(x)dx \approx \sum_{x_i \sim g(x)} f(x_i),$$

where x is a variable in $\mathbb{R}^n$, $g(x)$ is a probability distribution on it, while $f(x)$ is an arbitrary function of x . Reframing this equality, if we want to compute an integral in the form $\int f(x)dx$, we can do so by identifying an auxiliary density g , and then evaluating

$$\int f(x)dx = \int \frac{f(x)}{g(x)}g(x)dx \approx \sum_{x_i \sim g(x)} \frac{f(x_i)}{g(x_i)}. \quad (2.13)$$

The quality of this as an integration technique varies depending on the quality of the proposal g . In fact, it can be shown that the estimator's variance is related to the variance of $f(x_i)/g(x_i)$ [12]: the closer the proposal g approximates the integrand f (up to a constant multiplicative factor), the better.

This approach is applicable to gravitational wave inference if we are using NPE: we have direct access to the integrand function, $f = \mathcal{L}(d|\theta_i)\pi(\theta_i)$, and neural posterior estimation provides us with an estimate $q(\theta|d)$ of the posterior density, which we may use as in place of the proposal g , since its density is also available.

The procedure is then as follows: we extract n samples $\theta_i \sim q(\theta|d)$, and assign them weights

$$w_i = \frac{\mathcal{L}(d|\theta_i)\pi(\theta_i)}{q(\theta_i|d)}. \quad (2.14)$$

By Equation 2.13, then, the average of these weights will be an estimator for the evidence:

$$Z(d) \approx \frac{1}{n} \sum_{i=1}^n w_i.$$

While this is an unbiased estimator, its variance may be high if the proposal distribution is not close enough to the function we are integrating. Let us define the number of **effective samples** [12]:

$$n_{\text{eff}} = \frac{(\sum_i w_i)^2}{\sum_i w_i^2}, \quad (2.15)$$

and the **sample efficiency** $\epsilon = n_{\text{eff}}/n$.

i Note 3: Examples and intuition

The optimal situation is when the proposal distribution q perfectly matches the posterior $p(\theta|d)$: in that case, all the weights will be identical and equal to Z . Then, the effective sample size will be $n_{\text{eff}} = (nZ)^2/(nZ^2) = n$, and the efficiency ϵ will equal 1.

If the proposal poorly models the posterior, we have no strict mathematical guarantees on what the efficiency will be. In principle, depending on the likelihood's shape, it is possible for the weights to be roughly constant, despite none of the samples lying in the posterior's typical region.

In practice, though, gravitational wave likelihoods exhibit extreme variations far from their typical regions. The log-likelihood ratio in Equation 3.4 will range at the very least from $-\rho^2/2$ to $\rho^2/2$; even for low-SNR detections with $\rho \approx 10$, this is a variation of 100 in the *logarithm* of the likelihood. We observe that, when the proposal poorly models the posterior distribution, one weight tends to completely dominate the sum, so that (for instance) $w_1 \gg w_i$ for any $i \neq 1$. If this is the case, we get $n_{\text{eff}} \approx w_1^2/w_1^2 = 1$ and $\epsilon = 1/n$.

In [158, figure 1], one can see a scenario analogous to this occur: there, the number of samples drawn is $n = 10^5$; while for most sources the sample efficiency is high, for a few we can see very low values, with some reaching the lower bound of 10^{-5} , which corresponds to an effective sample size of approximately 1.

The evidence estimate can be associated with an approximate error [62]

$$\sigma_Z \approx Z \sqrt{\frac{1 - \epsilon}{n\epsilon}}$$

or equivalently

$$\sigma_{\log Z} \approx \sqrt{\frac{1 - \epsilon}{n\epsilon}}.$$

2.3.3.2. Numerical example

We will take the unnormalized posterior to be

$$\mathcal{L}(\theta|d)\pi(\theta) = 100\mathcal{N}(\theta, \mu = 0, \sigma = 1)$$

while the proposal distribution will be slightly offset from it:

$$q(\theta|d) = \mathcal{N}(\theta|\mu = \mu_0, \sigma = 1).$$

In order to illustrate the impact of effective sample size, we choose two different values for the offset: 2.5 and 0.3. In Figure 2.20 we can see the impact of effective sample size on the evidence estimate.

Besides the decreasing accuracy in evidence estimation, a low efficiency (low effective sample size) also leads to inaccurate estimates on the error of the estimate itself. Visually, this is not surprising: with the higher value of the offset, the left side of the posterior is completely unexplored.




Figure 2.20.: A demonstration of importance sampling with an efficient or an inefficient proposal.

2.3.3.3. Importance sampling for neural posterior estimation

Dax et al. [62] introduced the usage of importance sampling to complement neural posterior estimation, with a piece of software called DINGO-IS. The density estimate given by DINGO for the posterior is corrected through importance sampling, based on the standard gravitational wave likelihood.

Besides giving an estimate for the evidence, this allows us to correct the posterior samples. We can extract n samples θ_i from the NPE distribution $q(\theta|d)$, compute their weights according to Equation 2.14, and then interpret θ_i as *weighted* samples from the posterior distribution $p(\theta|d)$, and use them as described in Section 2.1.1.2.

This approach effectively makes the DINGO-IS analysis results equivalent to those obtained by stochastically sampling the likelihood. While bringing a significant advantage in terms of robustness and trustworthiness, this approach — if applied with a Gaussian likelihood — negates any benefits that could be gained from training the neural network on anything other than Gaussian noise.

Analyzing a signal with DINGO-IS can be both significantly faster and significantly slower than performing the same analysis with stochastic methods such as nested sampling: when the sample efficiency is low, the number of samples required for unbiased inference can quickly become impractical.

2.4. Fisher matrix

The Fisher matrix is a common tool in gravitational wave data analysis and forecasting. It is defined as the noise-averaged Hessian of the log-likelihood:

$$F_{ij} = \langle \partial_i \partial_j \log \mathcal{L}(d|\theta) \rangle_{\text{noise}}.$$

If the likelihood is written in the Whittle form (see Section 3.2), then this is equal to

$$F_{ij} = \left(\frac{\partial h}{\partial \theta_i} \middle| \frac{\partial h}{\partial \theta_j} \right).$$

For a derivation, see Finn [76] or Vallisneri [186]. This matrix can be interpreted in different ways [186]: *(i)* in terms of the frequentist Cramér-Rao bound on the error covariance of an unbiased estimator for the true source parameters, *(ii)* as the error covariance of the frequentist maximum-likelihood estimator, or *(iii)* as an estimate of the error covariance of the Bayesian posterior distribution under the assumption that the priors are constant over the range of interest; here we shall focus on option *(iii)* for consistency with the Bayesian approach in the rest of this work.

We can approximate the likelihood in the linearized signal approximation (LSA), a Taylor series for the signal around the true value of the parameters θ_0 in terms of the deviation $\delta\theta = \theta - \theta_0$:

$$h(\theta) = h(\theta_0) + \delta\theta_i \frac{\partial h}{\partial \theta_i} + \mathcal{O}((\delta\theta)^2).$$

In this approximation, the Whittle likelihood is written as

$$\log \mathcal{L}(d|\theta) = -\frac{1}{2}(n|n) + \delta\theta_i \left(n \middle| \frac{\partial h}{\partial \theta_i} \right) - \frac{1}{2} \delta\theta^i \delta\theta^j \mathcal{F}_{ij} + \mathcal{O}((\delta\theta)^3). \quad (2.16)$$

This is the expression of a multivariate normal distribution, with covariance matrix $\mathcal{C}_{ij} = \mathcal{F}_{ij}^{-1}$.

This expression is written as an expansion in orders of the deviation of the parameters, but the actual requirement for the Fisher approximation is a high value of the (optimal) SNR, which we shall define in Section 3.3, $\rho = \sqrt{(h_0|h_0)}$. In this approach, and based on the LSA approximation to the likelihood, we can compute the first and second moment of the posterior distribution $p(\theta|d) \propto \mathcal{L}(d|\theta)\pi(\theta)$ as a series expansion in orders of $1/\rho$. This is only analytically tractable if the prior distribution $\pi(\theta)$ is constant over the range of interest, or Gaussian and centered around the same point as the posterior.

The expectation value of the noise covariance is then

$$\langle \delta\theta_i \delta\theta_j \rangle_{p(\theta|d)} = F_{ij}^{-1} + \mathcal{O}\left(\frac{1}{\rho^3}\right) ;$$

with the addition of a Gaussian prior $\pi(\delta\theta) \propto \exp\left(-\frac{1}{2}P_{ij}\delta\theta_i\delta\theta_j\right)$ it becomes [186, eq. 39]:

$$\langle \delta\theta_i \delta\theta_j \rangle_{p(\theta|d)} = (F_{ij} + P_{ij})^{-1} + \mathcal{O}\left(\frac{1}{\rho^3}\right) .$$

The Fisher matrix can be complemented by the addition of information from a generic prior distribution by performing rejection sampling on samples extracted from the Gaussian distribution it defines [95, 67].

Several studies [186, 150, 67, 159], some of which include prior distributions through rejection sampling, have found that the covariance matrix estimated by inverting the Fisher matrix does not reliably represent the posterior covariance matrix obtained through the sampling techniques described in this chapter.

Nevertheless, the Fisher matrix can be a useful tool, especially in the context of forecasting, due to the relative ease of its computation compared to full parameter estimation. It provides a metric on the parameter space, which can be useful when generating template banks, see e.g. Phukon, Schmidt, and Pratten [141]. In some cases it can be computed analytically, which is useful for the theoretical understanding of the behavior of data analysis pipelines, see Section 2.4.1.

Fisher integrands are also a useful tool for the understanding of parameter constraints: a quantity in the form

$$\frac{4f}{S_n(f)} \left| \frac{\partial f}{\partial \theta_i} \right|^2$$

can be integrated in $d \log f$ to yield the i -th entry in the diagonal of the Fisher matrix, whose inverse contributes to the computation of the variance of the i -th parameter. It is

incomplete information, since it does not account for correlations with other parameters; still, it can provide some qualitative understanding and diagnostics. We used this tool to describe the way information on the sky position of the source accumulates as a function of time with the Lunar Gravitational Wave Antenna [56].

In practice, the computation of Fisher matrices requires differentiation of the projected waveform onto a detector. This can be accomplished numerically, but it is generally convenient to analytically compute all derivatives where it is possible [68]. For example, the derivative with respect to distance is simply

$$\frac{\partial h}{\partial d_L} = -\frac{h}{d_L},$$

since $h \propto 1/d_L$. If the waveforms and projection allow it, even more of the derivatives can be computed analytically, with the help of automatic differentiation such as that enabled by jax [95, 97].

2.4.1. Analytic Fisher calculations

2.4.1.1. Timing uncertainty

We can obtain a Fisher estimate for the timing uncertainty of a gravitational wave observation by computing the Shur’s complement of the Fisher matrix for time and phase [87, 163].

The result is expressed in terms of the frequency moments of the signal:⁹

$$\langle f^k \rangle = \frac{1}{\rho^2} (f^k h | h).$$

With these, we can define the frequency bandwidth

$$\sigma_f^2 = \langle f^2 \rangle - \langle f \rangle^2$$

and from it the timing uncertainty

$$\sigma_t = \frac{1}{2\pi\rho\sigma_f}.$$

⁹In some references [73, 72] this expression is incorrectly reported as $\langle f^k \rangle = (f^k h | h)$ (or equivalently in each paper’s notation). Without the normalization factor ρ^2 , the two terms in the frequency bandwidth expression scale with ρ^2 and ρ^4 respectively, meaning that for high enough SNR the variance becomes negative.

2.4.1.2. Angular resolution

Localizing gravitational wave sources is crucial, and it is quite useful to analytically understand which effects drive it in order to inform detector design. Wen and Chen [192] computed a general Fisher expression for the angular area of a source observed by a network of detectors.

This area is a solid angle defined as

$$\Delta\Omega_{1\sigma} = 2\pi\sqrt{\langle\delta n_x^2\rangle\langle\delta n_y^2\rangle - \langle\delta n_x\delta n_y\rangle^2} = 2\pi\sqrt{\det C_{\text{sky}}},$$

where C_{sky} is the two-dimensional covariance matrix for the two components of the sky position position vector, n_x and n_y , which are orthogonal to its true value n_z . In a spherical coordinate system with right ascension α and declination δ we need to add a factor $|\cos \delta|$ to get the correct area element [95]. This area is defined such that the the probability contained within a similarly shaped contour of area $\Delta\Omega$ around the center is $P(\Delta\Omega) = \exp(-\Delta\Omega/\Delta\Omega_{1\sigma})$, therefore the equivalent area at the X credible level is

$$\Delta\Omega_X = \Delta\Omega_{1\sigma} \log(1 - X).$$

The Fisher expression for the 1σ area reads [48]:

$$\Delta\Omega_{1\sigma} = \frac{4\sqrt{2}\pi c^2}{\sqrt{\sum_{JKLM} \Delta_{JK}\Delta_{LM} |(\mathbf{r}_{KJ} \times \mathbf{r}_{ML}) \cdot \mathbf{n}|^2}},$$

where the indices $JKLM$ span the available detectors, Ω is the angular frequency, whereas

$$\Delta_{IJ} = \langle\Omega d_I|\Omega d_J\rangle\delta_{IJ} - \langle\Omega d_I|\partial_k d_J\rangle\langle\partial_l d_J|\Omega d_J\rangle F_{kl}^{-1}$$

where F_{kl} is the Fisher matrix on all the parameters describing the signal except for sky position.

The term $|(\mathbf{r}_{KJ} \times \mathbf{r}_{ML}) \cdot \mathbf{n}|$ is twice the area formed by the projections of the detectors J, K, L, M onto the plane orthogonal to the wave propagation direction.

This expression is quite unwieldy in the general case; however, it can be simplified significantly for signals whose intrinsic parameters are assumed to be known.

For a two-detector network observing a monochromatic wave of (known) frequency f_0 , we get

$$\Delta\Omega_{2\text{ det}} = \frac{1}{\rho_N} \frac{c}{\pi f_0 D_\perp} \sqrt{\frac{\rho_N^4}{\rho_1^2 \rho_2^2}},$$

while for three detectors we get

$$\Delta\Omega_{3\text{ det}} = \frac{1}{\rho_N^2} \frac{c^2}{4\pi f_0^2 A_\perp} \sqrt{\frac{\rho_N^6}{\rho_1^2 \rho_2^2 \rho_3^2}}.$$

Here, $\rho_N = \sqrt{\sum_i \rho_i^2}$ is the network SNR, the sum in quadrature of the individual SNRs (see Section 3.3); the quantities $D_\perp$ and $A_\perp$ are respectively the distance between detectors or their area, projected onto a plane orthogonal to the propagation direction n .

An expression can also be computed in the case of a long observation by a single detector, since its orbit around the Sun will effectively make it akin to a large network of detectors. Wen and Chen [192] give this expression only in the limit of a short observation or one that is longer than a year; the expression for an arbitrary-length observation of a monochromatic signal with frequency f_0 reads [183]

$$\Delta\Omega = \frac{c^2}{f_0^2 \rho^2 A_s}, \quad (2.17)$$

where A_s is the area of the projected detector's motion onto a plane perpendicular to the source direction.

Specifically, it is defined as

$$A_s = 2\pi \sqrt{\text{var}(r_\theta) \text{var}(r_\phi) - \text{cov}^2(r_\theta, r_\phi)}, \quad (2.18)$$

where the coordinates r_θ and r_ϕ parameterize the projected motion of the detector in the aforementioned plane. When computing these variances and covariances, the trajectory should be weighted according to the accumulation of SNR, with an expression in the form

$$w(t) = \left| \frac{df}{dt} \right| \frac{|h(f)|^2}{S_n(f)}.$$

This is a consequence of equation 45 in the paper by Wen and Chen [192], under the assumption of a monochromatic signal.

2.4.1.3. Correlation structure

The Fisher matrix can aid in understanding the correlation structure of the posterior distribution.

For example, Singer and Price [163] used it to come to the conclusion that the extrinsic parameters of a gravitational wave source are overall uncorrelated from the intrinsic ones (for definitions see Chapter 5 and Chapter 6). This is qualitatively due to the fact that the extrinsic parameters are generally measured by *inter-detector* amplitude and time of arrival *differences*, while intrinsic parameters are generally measured by *averages* of amplitude and phase across detectors.

We can validate this by computing the Pearson correlation matrix of a full parameter estimation run:

$$\rho_{ij} = \frac{C_{ij}}{\sigma_i \sigma_j},$$

where C_{ij} is the empirical covariance matrix computed from the posterior samples, while $\sigma_i = \sqrt{C_{ii}}$ is the standard deviation of parameter i .

Each entry ρ_{ij} will lie between -1 and 1, with values near these extremes representing near-perfect (anti)correlation, while values near 0 represent uncorrelated variables. This only captures *linear* correlations, so it does not tell the whole story, but it is still quite useful.

In Figure 2.21 we see the structure of the correlation matrix for an actual analysis of the neutron star binary GW170817 (specifically, the one from Tissino et al. [182]).




Figure 2.21.: Correlations between some measured parameters for GW170817. c_{AL1} and $c_{\phi L1}$ are two calibration parameters for the Livingston detector, included for reference, out of the ~ 12 calibration parameters per detector used in real analyses. Arrival time and overall phase are not included since they were analytically marginalized in this particular analysis.

3. Gravitational wave parameter estimation

In this chapter, we outline the basic tools used for the estimation of gravitational wave signal parameters, including the Whittle likelihood (Section 3.2), the full time-domain likelihood (Section 3.2.2), strain visualization tools (Section 3.2.3) and likelihood acceleration techniques (Section 3.4).

Furthermore, in Section 3.2.1 we summarize the main result of a recent work on the (lack of) necessity to correct the likelihood for the loss of signal power incurred by windowing the data.

3.1. Fourier transform basics

The starting point for our analysis is typically an estimate for the gravitational wave strain observed by one or more detectors, which is expressed as a discretized time series $d(i\Delta t)$, with integer i . Each of the entries in this time series is noisy, and they cannot be considered in isolation: due to the wide-band nature of the noise, any given pair of points $d(i\Delta t)$ and $d(j\Delta t)$ will in general be *correlated*. Then, in order to analyze a data segment of duration T , with $N = T/\Delta t$ samples, we would naïvely need to consider all N^2 combinations.

The Fourier transform of the signal can come to our aid: its continuous version is

$$\tilde{h}(f) = \int_{-\infty}^{\infty} h(t) \exp(-2\pi i f t) dt,$$

while the discrete version, computed at a frequency $f = m/T$ with integer m , is

$$\tilde{h}(m/T) = \Delta t \sum_{n=0}^{N-1} h(n\Delta t) \exp\left(-\frac{2\pi i m n}{N}\right). \quad (3.1)$$

i Note 4: Fourier transform periodicity and range

While Equation 3.1 can be computed for any $m \in \mathbb{Z}$, it only contains independent information for $0 \leq m < N/2$: since our time domain signal is real-valued we have

$\tilde{h}(-m/T) = \tilde{h}^*(m/T)$, and by the periodicity of the complex exponential $\tilde{h}(m/T) = \tilde{h}((m+N)/T)$.

Starting from N real-valued samples in time we then get $N/2$ complex-valued ones in frequency. The frequency samples are spaced by $1/T$, while the highest frequency available is $1/2\Delta t = f_s$, the Nyquist frequency. Any component in the signal with a higher frequency, $f > f_s$, will experience aliasing: it will appear in the Fourier transform of the data at an artificially lower frequency, $f - kf_s$ for some integer k chosen such that $0 < f - kf_s < f_s$.

The Fast Fourier Transform algorithm allows us to compute this relatively quickly, in $O(N \log N)$ operations [91].

Under some assumptions, the Fourier components of the noise are uncorrelated among each other. Because of this, we will start our discussion with the Fourier-domain Whittle likelihood (Section 3.2), which has an $O(N)$ complexity, but also describe the way parameter estimation can be performed when those assumptions break down and the Fourier-domain noise is not diagonal (Section 3.2.2).

3.2. The Whittle likelihood

The likelihood expresses the answer to the question:

How likely are we to have observed the data d , under the assumption that this data can be decomposed as noise n plus a signal $h(\theta)$ with fixed parameters θ ?

If we are writing the data as $d = n + h(\theta)$, this is equivalent to asking:

How likely are we to observe a noise realization equal to $n = d - h(\theta)$?

In order to answer this, we need a model for the noise. In gravitational wave data analysis, it is commonly assumed that the noise can be described as a stationary time series, $n(t)$, whose autocorrelation $\mathbb{E}[n(t)n(t + \Delta t)] \equiv \gamma(\Delta t)$ can be expressed as a function of Δt only [70, section II.C]. A further assumption made is that the noise has zero mean ($\mathbb{E}[n(t)] = 0$ for any time).

The Fourier transform of the autocorrelation function $\gamma(\Delta t)$ is called the Power Spectral Density, $S_n(f)$, which can alternatively be expressed as [125]:

$$\langle n(f)n^*(f') \rangle = \frac{1}{2} \delta(f - f') S_n(f).$$

As defined, this is the *one-sided* PSD, which satisfies

$$\langle |n(t)|^2 \rangle = \int_0^\infty S_n(f) df.$$

We could have alternatively defined it without the $1/2$ factor, simultaneously extending the integral from $-\infty$ to ∞ . The negative-frequency contribution to the integral is equal to the positive-frequency one, since the signal and noise are real-valued. Sometimes, the *Amplitude Spectral Density* (ASD) is also quoted, which in gravitational wave data analysis is defined as the square root of the PSD, with units of $\text{Hz}^{-1/2}$:¹

$$A_n(f) = \sqrt{S_n(f)}. \quad (3.2)$$

The PSD can be used for the characterization of the second moment of any stationary noise; if we further assume that the noise is Gaussian and the data are periodic,² we can obtain the Whittle approximation to the Gaussian likelihood [193, 154]:

$$\begin{aligned} \log \mathcal{L}(n) &= \log \mathcal{N}_1 + \sum_{m>0} \left[\log \left(\frac{2}{\pi T S_n(f_m)} \right) - \frac{2|n(f_m)|^2}{T S_n(f_m)} \right] \\ &= \log \mathcal{N}_2 - \frac{1}{2} \int_0^{f_{\max}} \frac{4|n(f)|^2}{S_n(f)} df \\ &= \log \mathcal{N}_2 - \frac{1}{2} (n|n), \end{aligned}$$

where T denotes the duration of the data segment.

The normalization terms $\mathcal{N}_1$ and $\mathcal{N}_2$ are different, but crucially they only depend on the Power Spectral Density, not on the specific noise realization.

In the last equation, we introduced the notation $(\cdot|\cdot)$ for the noise-weighted inner product between waveforms, defined as

$$(a|b) = 4\Re \int_0^\infty \frac{a(f)b^*(f)}{S_n(f)} df, \quad (3.3)$$

where z^* is the complex conjugate of z , while $\Re$ denotes the real part operator, defined so that $\Re z = (z + z^*)/2$.

More specifically, this likelihood answers the question:

¹More in general, the ASD has units of $x/\sqrt{\text{Hz}}$, where x is the unit of the time-domain noise. For example, if we were characterizing acceleration noise x would be m/s^2 .

²For a detailed discussion of the periodicity assumption, see Section 3.2.1.

How likely is it that noise distributed according to the PSD $S_n(f)$ will have the specific realization n ?

It is possible to infer the PSD together with the signal parameters [70], but in the data analysis for gravitational wave signals from ground-based interferometers, it is common practice to fix the PSD to an estimate obtained considering data near the event. This is acceptable, since variations in the noise distribution typically happen on much slower time scales than duration of the signal itself.

In such a case, we can rewrite the likelihood in terms of the signal $h(\theta) = d - n$ as follows:

$$\begin{aligned}\log \mathcal{L}(d|\theta) &= \log \mathcal{N}_2 - \frac{1}{2}(d - h(\theta)|d - h(\theta)) \\ &= \log \mathcal{N}_2 - \frac{1}{2}(d|d) + \Re(h(\theta)|d) - \frac{1}{2}(h(\theta)|h(\theta)).\end{aligned}$$

We can work around the messy normalization term [154]. Our likelihood is written assuming that a signal is present in the data, but supposing that was not the case, we would have $d = n$ and therefore

$$\log \mathcal{L}(d|\text{noise only}) = \log \mathcal{N}_2 - \frac{1}{2}(d|d).$$

Since we are assuming the PSD is fixed, this model has zero parameters, and this likelihood can also be interpreted as an evidence: $\log \mathcal{L}(d|\text{noise only}) = Z_{\text{noise}}$.

Therefore, we can perform parameter estimation using the likelihood ratio

$$\log \mathcal{L}_r(d|\theta) = \log \frac{\mathcal{L}(d|h(\theta))}{Z_{\text{noise}}} = \Re(h(\theta)|d) - \frac{1}{2}(h(\theta)|h(\theta)), \quad (3.4)$$

and the evidence we will obtain, the integral of $\mathcal{L}_r(d|\theta)\pi(\theta)d\theta$ over the parameter space, will actually be a Bayes Factor, $Z_{\text{signal and model}}/Z_{\text{noise}}$.

3.2.1. Non-periodic data: windowing

The Fourier transform in equation Equation 3.1 is, technically speaking, only appropriate if the data is *periodic*. This is not the case for realistic data: on the contrary, we only want to consider a given segment in time around our event. Then, the noise matrix will not be exactly diagonalized by the Fourier transform.

We can model the process of selecting a finite chunk of data from a theoretically infinite process as multiplication by a window function $w(t)$ [174]: if the data time series is d_i , originating from an infinite process D_i , we can define a corresponding window time series w_i and get

$$d_i = D_i w_i .$$

The series w_i is here simply defined to be 1 for indices inside the window, and 0 outside. In the Fourier domain, multiplication is replaced by convolution:

$$\tilde{d}_k = (\tilde{D} * \tilde{w})_k = \Delta t \sum_{j=0}^{N-1} D_j w_j \exp\left(-\frac{2\pi i j k}{N}\right) .$$

Any window function will therefore introduce *spectral leakage*, making the noise covariance significantly non-diagonal in the frequency domain.

In gravitational wave science, the Tukey window is commonly chosen:

$$w(t) = \begin{cases} \frac{1}{2} [1 - \cos(\frac{2\pi t}{\alpha T})] & 0 \leq t \leq \alpha T/2 \\ 1 & \alpha T/2 \leq t \leq T/2 \\ w(T-t) & 0 \leq t \leq T \end{cases}$$

It significantly decreases spectral leakage compared to the rectangular window, as seen in Figure 3.1, despite leaving the central $1 - \alpha$ seconds of the data unchanged.

Applying a Tukey function with parameter α to a stationary time series will, on average, reduce its power to a fraction of the total given by [173]

$$\beta = \int w^2(t) dt = 1 - \frac{5\alpha}{8} .$$

Before mid-2025, it was common practice in the analyses performed by the LIGO-Virgo-KAGRA collaboration to counteract this effect by rescaling the whole data segment by a factor $1/\sqrt{\beta}$.

However, Talbot et al. [173] showed that

1. the non-diagonality in the noise matrix introduced by the Tukey window has no significant impact on inference if the data segment is long enough;³

³The precise requirement is as follows: let us define D_w as the diagonal matrix which applies the whitening operation in the time domain, and the whitening filter $\mathcal{W}$, written in terms of the noise covariance matrix C as $\mathcal{W}^\top \mathcal{W} = C^{-1}$. Then, the impact on inference is small as long as the two quantities $\mathcal{W}h$ and $\mathcal{W}(1 - D_w)d$ do not significantly intersect.




Figure 3.1.: A comparison of a rectangular and Tukey window function in the time and frequency domains.


2. rescaling the data by $1/\sqrt{\beta}$ brings about a small, but detectable, systematic bias in inference.

The second point was shown through a pp-plot. For a few thousand times, a signal was injected and recovered in a data segment, with or without the correction. The recovery was performed only varying the luminosity distance of the signals.

It shows signs of overconstraining: compare with the example in Figure 2.7. Rescaling the data is equivalent to multiplying the log-likelihood by $1/\beta$, which makes its peak artificially too narrow.

3.2.2. Non-periodic data: time domain parameter estimation

The procedure of windowing and tapering the data works well enough for the analysis of a full gravitational wave signal, for which we can consider a relatively long segment around the event. However, in some cases we do not have that luxury: for example, the ringdown is the last section of the signal from a black hole coalescence, in which the remnant black hole settles to a quiescent state (see Section 4.2). In the analysis of this part of the data, it is crucial to be able to isolate the perturbative regime, which only describes the data well starting a certain time after the merger.



(a) p - p plot for the distance-only injections with and without the correction. (b) K-S test p -value as a function of the number of injected signals. The p -value is shown on a logit scale (Note 1).

Figure 3.2.: Tests of the impact of the window factor on parameter estimation, adapted from Talbot et al. [173].

To this end, we want to only analyze data starting from a specific point in time which “cuts off” part of the signal in the region where it is stronger. Trying to impose this kind of sharp time-domain cut in the frequency domain is a losing endeavour. The solution is to go back to a time-domain likelihood, in the form:

$$\mathcal{L}(d|\theta) = -\frac{1}{2} \sum_{i,j=0}^N (d_i - h_i) C_{ij}^{-1} (d_j - h_j) + \text{const.}$$

where d_i and h_i are the discretized time-domain data and predicted strain, defined so that the index $i = 0$ corresponds to the cut time [98].

The time-domain covariance C_{ij} is not diagonal, unlike the frequency-domain one. Still, there are some computational tricks one can use.

We need two definitions: a matrix is said to be *Toeplitz* if its diagonals contain constant values: for all $0 \leq i, j < N$, we have $C_{ij} = C_{(i+1)(j+1)}$. An $N \times N$ Toeplitz matrix has $2N - 1$ independent entries. A *symmetric Toeplitz* matrix, furthermore, satisfies $C_{ij} = C_{ji}$, and only has N independent entries. It can be expressed in terms of a vector ρ with N entries: $C_{ij} = \rho_{|i-j|}$. A *circulant* matrix is a Toeplitz matrix in which every row is equal to the previous one, shifted by one: $C_{ij} = \rho_{(i-j) \bmod N}$; if the circulant matrix is symmetric, then $C_{ij} = \rho_{|i-j| \bmod N}$. A symmetric circulant metric only has $\lceil N/2 \rceil$ independent entries.

The time-domain noise covariance of a stationary process is always symmetric Toeplitz: its entry C_{ij} describes the covariance between time $i\Delta t$ and time $j\Delta t$; by stationarity, this can only depend on the delay $(i - j)\Delta t$, therefore only on the difference $i - j$. As this

is a covariance matrix it will be symmetric: the entries can actually only depend on the modulus of the index difference, $|i - j|$. Furthermore, if we impose periodic boundary conditions then the covariance is also circulant; we cannot do so in the time-limited case.

Circulant matrices are exactly diagonalized by the discrete Fourier transform. Toeplitz matrices are not, but they can be fully characterized by the autocovariance function ρ : $C_{ij} = \rho(|i - j|)$.

The time-domain likelihood can be written as in Equation 3.4, but with scalar products in the form [98]

$$(a|b)_{\text{TD}} = \sum_{i,j=0}^{N-1} a_i C_{ij}^{-1} b_j.$$

If C is Toeplitz, this can be efficiently computed with a standard routine such as Levinson–Durbin recursion, which returns a the solution x to a linear system $b_i = C_{ij}x_j$ when C is a Toeplitz matrix. Alternatively, one can compute the lower-triangular Cholesky factor L such that $LL^\top = C$, invert it, and then compute the product as the scalar product of the *whitened* vectors $L^{-1}a$ and $L^{-1}b$

$$(a|b)_{\text{TD}} = (L^{-1}a) \cdot (L^{-1}b).$$

3.2.3. Strain plotting and whitening

It is common for data analysts to wish to start their description of some signal by showing the corresponding *unprocessed* data, the starting point for their endeavors to draw physical meaning. Being able to show unprocessed data is also useful in the context of scientific communication, since it allows the presenter to show a figure which can be immediately understood.

This task is made tricky in the context of gravitational wave analysis by the nature of the noise. The sensitivity “bucket” in the frequency domain means that the magnitude of the noise outside a given band completely overwhelms the time-domain data: in ground-based observations, while typical strain magnitudes are on the order of 10^{-21} or less, raw time-domain strain data is often $> 10^{-18}$.

If we are determined to show strain, we need to filter the series: Figure 3.3 shows this for GW250114. A Butterworth band-pass filter between 35Hz and 350Hz was applied twice to the data, in addition to a notch filter at 60Hz to remove power grid-driven noise. The original detection paper for GW150914 [114] used many more hand-tuned notches to eliminate the numerous spectral features present in the noise at that time, which have fortunately been reduced nowadays.




Figure 3.3.: GW250114, shown in the same way as the first figure from the discovery paper for GW150914, LIGO Scientific Collaboration and Virgo Collaboration et al. [114].

Besides the apparent technical complexity of this approach and its many parameters, its residuals are problematic. Residuals computed subtracting a filtered model from filtered strain data have little meaning, since their distribution is hard to model and frequency-dependent: one should refrain from trying to draw visual conclusions based on their patterns and magnitude, which essentially removes all of the potential utility in showing them. Furthermore, the filtering process can introduce artefacts in the data: the fact that the signal model does not exactly approach zero after the merger is one of these.

A significant advantage of showing filtered data, on the other hand, is being able to show the actual signal magnitude, on the order of 10^{-21} , on the y -axis. The alternative approach, which has become the standard nowadays, does not share this.

Whitening the data means normalizing it to the noise. Applying the whitening procedure to a noise-only data stream returns white noise. In the frequency domain, whitening a series $h(f)$ means defining

$$h_w(f) = h(f) \sqrt{\frac{4}{TS_n(f)}},$$

where S_n is the noise power spectral density, while T is the segment duration. With this normalization, each noise entry has variance 2 (due to its independent real and imaginary components) [24]. This can also be generalized to time-domain analyses, as discussed in Section 3.2.2.

In the presence of a signal, whitened time-domain strain will deviate from its mean of zero; the deviation amount can be immediately interpreted as a number of standard deviations. For GW250114, this reached the level of 15σ , as one can see in Figure 4.2.

3.3. Matched filtering and the signal-to-noise ratio

The discussion so far has discussed the analysis of a signal we know to be present in the data. However, the tools we described can also be used to quantify the *detectability* of a signal.

The *matched-filter* signal-to-noise ratio (SNR) statistic ρ for the data d , assuming a template h , is given by

$$\rho_{\text{MF}} = \frac{(d|h)}{\sqrt{(h|h)}}$$

in terms of the scalar products defined in Equation 3.3.

A common approach, when looking for a signal in the data, is to look for high values in this quantity for a large set of templates h .

If we fix the template h we can also compute the expectation value of the SNR when varying the noise realization:

$$\rho_{\text{opt}} = \mathbb{E}_n \left[\frac{(d|h)}{\sqrt{(h|h)}} \right] = \sqrt{(h|h)}. \quad (3.5)$$

This is called *optimal* SNR, though it involves no optimization procedure. If a template with a given optimal SNR is injected into Gaussian noise, the matched-filter SNR is equally likely to be higher and lower than the optimal value.

Since the optimal SNR depends on an integral in the form $\rho_{\text{opt}}^2 \sim \int |h|^2 / S df$, independent SNRs (such as ones from different detectors, or different frequency bands considered independently) add in quadrature:

$$\rho_{\text{tot}}^2 = \sum_i \rho_i^2.$$

The highest SNR for a gravitational wave signal, as of early 2026, was achieved by the event GW250114 [2], which reached a value between 77 and 80 when combining the Livingston and Hanford LIGO detectors.

3.3.1. SNR thresholds

This section gives a brief outline of the way to estimate what value of SNR is sufficient for a signal to be “detectable”. This is useful information in the context of forecasting the number of detected signals by some detector that is not yet operational.

If we were searching for a signal without any possible shift in time and with the certainty that the noise did not contain any glitches or non-Gaussianities, then the SNR would simply follow a χ^2 distribution with two degrees of freedom [13], meaning that we could reach “five σ ” significance (p -value of 5.7×10^{-7}) with a threshold of $\rho = 5.5$.

This is not the case in real data, since both the aforementioned assumptions are false: the different possible arrival times of a signal are effectively “more trials”, but more importantly the non-Gaussianity of the noise means that high-SNR triggers in a matched filtering search will be much more common than simple Gaussian statistics would indicate. Accounting for these factors is tricky, and in real data it is typically accomplished by performing time-slides: taking real detector data, but shifted across detectors by more than $t_{\text{min}} = d/c$, where d is the maximum distance separating them. In the assumption that signals are rare (which holds for current detectors, for ET this will have to be challenged), this allows us to generate arbitrary amounts of noise-only realistic data by removing all possible coincidences. We can then compute the rate of noise instances surpassing our threshold, even if this rate

is very low. With this, we can compute the false-alarm rate (FAR) associated to a given SNR threshold. It should be noted that real searches are significantly more sophisticated than SNR maximizers; for example, PyCBC requires that each frequency band contributes appropriately to the SNR [132], which allows us to discard many loud glitches.

Nevertheless, high SNR is correlated with low FAR for real signals, and the scaling of the FAR as a function of SNR, in the absence of large outliers, is typically exponential [118]:

$$\text{FAR} \approx \text{FAR}_8 \exp\left(-\frac{\rho - 8}{\alpha}\right).$$

The specific values for the FAR at $\rho = 8$ and the scaling parameter α are found through a fit, and they depend on the signal template. A fit of O1 data found $\alpha = 0.18$ and $\text{FAR}_8 = 5500\text{yr}^{-1}$ for BBH, and $\alpha = 0.13$ and $\text{FAR}_8 = 30000\text{yr}^{-1}$ for BNS. This value of α means that the exponential cutoff is quite steep: for example, bringing the BNS FAR from 30 thousand per year to just one per year only requires us to raise the equivalent threshold to $\rho = 9.4$ on average.

Based on this, we can determine the SNR threshold we should consider as a minimum for detectability in forecasting. Despite not knowing which glitches will affect our future detectors, this sharp exponential behavior gives us confidence that values in the range of roughly 8 to 12 will be sufficient to guarantee that the number of false positive detections is low enough.

What constitutes a “low enough” rate of false positives is subjective, and the choice of a threshold will depend on the rate of true detections for a given type of source, as well as other considerations. For example, in the fourth observing run of the LIGO-Virgo-KAGRA collaboration (O4), alerts were labelled “significant” with a false alarm rate of less than one per month when relative to compact binaries, but they required a FAR of less than one per year if they were relative to bursts [45].

3.3.2. Noise plotting and characteristic strain

A common task in gravitational wave analysis and forecasting is to compare signal and noise for a given observation as a function of frequency. Qualitatively, most detectors’ noise PSDs $S_n(f)$ exhibit a sensitivity “bucket”, a frequency range in which they are most sensitive, with noise rising to either side. We would like to compare this bucket to the signal’s Fourier spectrum, $\tilde{h}(f)$. Unfortunately, simply showing them on the same axis does not make sense, even though they have the same unit of Hz^{-1} . To understand this, consider the equivalent scenario if we were comparing the PSD and a Fourier spectrum for displacement: then, the PSD would have units of m^2/Hz , while the Fourier domain signal would have units of m/Hz . This discrepancy is hidden by the fact that our signal,

strain, is dimensionless, but it remains a conceptual problem, and qualitatively the orders of magnitude of these two quantities make it very obvious.

In order to show these quantities together in a useful manner, we can look at the integral which defines the optimal SNR in Equation 3.5:

$$\rho_{\text{opt}}^2 = 4 \int \frac{|\tilde{h}(f)|^2}{S_n(f)} df.$$

The ratio in the integrand is the relevant comparison to determine the clarity of a signal with Fourier transform $|\tilde{h}(f)|$ measured by a detector with a PSD $S_n(f)$. This ratio still has units of Hz^{-1} , and we generally want to show this quantity as a function of logarithmic frequency: both of these issues are solved by rewriting the integral as

$$\rho_{\text{opt}}^2 = \int \left(\frac{2f|\tilde{h}(f)|}{\sqrt{fS_n(f)}} \right)^2 d \log f = \int \frac{h_c^2(f)}{h_n^2(f)} d \log f.$$

This defines the *characteristic signal strain* $h_c(f) = 2f|\tilde{h}(f)|$ and the *characteristic noise strain* $h_n(f) = \sqrt{fS_n(f)}$, both of which are dimensionless.

For more details on this and other ways to compare signal and noise, including for example computing the energy density contained in gravitational waves, see Moore, Cole, and Berry [125].

3.4. Likelihood acceleration

This section will discuss common approaches which can be used to accelerate the computation of a frequency-domain Whittle likelihood.

The two terms we need to compute for the likelihood ratio (Equation 3.4) are the following integrals:

$$\Re(h(\theta)|d) = 4\Re \int_0^{f_{\max}} \frac{h(f; \theta)d^*(f)}{S_n(f)} df,$$

and

$$(h(\theta)|h(\theta)) = 4 \int_0^{f_{\max}} \frac{|h(f; \theta)|^2}{S_n(f)} df.$$

The baseline approach is to evaluate these as Riemann sums on the full Fourier-domain grid (also called the “FFT grid”), which as we outlined in Note 4 has $T/2\Delta t$ frequency bins. This becomes less and less practical the longer the data segment is.

Several approaches to accelerate this computation have been proposed, see section 3 in Roulet and Venumadhav [154] for a review. Here we will outline two of the most popular: relative binning and reduced order quadrature.

The efficacy of both approaches is fundamentally driven by their ability to decrease the number of frequencies at which the waveform approximant is evaluated. This will generally reduce the total evaluation time, but only to a point: in Figure 5.11 we illustrate that the general scaling of waveform model evaluation takes the form $t = t_1 + Nt_2$, where N is the number of grid points. Decreasing N significantly below t_1/t_2 is going to exhibit diminishing returns.

3.4.1. Relative binning

Relative binning [196] is a technique to accelerate the evaluation of the integrals in the likelihood using the key idea that, while the waveform itself may be highly oscillatory, all the waveforms which are compatible with the data will be quite similar to each other (in a specific sense).

Starting from a reference waveform $h_0(f)$, we define the ratio $r(f) = h(f)/h_0(f)$ of the generic waveform h to the reference one h_0 . A good reference waveform can be computed, for example, by looking for the maximum value of the likelihood. This is still a “cheap” operation compared to sampling, and can be performed at full resolution.

Then, we can define a relatively coarse grid in frequency, and approximate $r(f)$ to linear order within each bin b : for all frequencies f in the bin,

$$r(f) = r_0(b) + r_1(b)(f - f_m(b)), \quad (3.6)$$

where $f_m(b)$ is the central frequency in the bin.

Then, we compute some summary statistics: for each bin b , define

$$\begin{aligned} A_0(b) &= \int_{f \in b} \frac{d(f)h_0^*(f)}{S_n(f)/T} \\ A_1(b) &= \int_{f \in b} \frac{d(f)h_0^*(f)}{S_n(f)/T} (f - f_m(b)) \\ B_0(b) &= \int_{f \in b} \frac{|h_0(f)|^2}{S_n(f)/T} \\ B_1(b) &= \int_{f \in b} \frac{|h_0(f)|^2}{S_n(f)/T} (f - f_m(b)). \end{aligned}$$

These integrals need to be computed at maximum resolution in frequency, but only once, with the reference waveform. With these at hand, the two terms are approximated as follows, in terms of a set of frequencies $b = \{f_i\}$:

$$\mathcal{L}_{hd} = \Re \sum_b (A_0(b)r_0^*(b) + A_1(b)r_1^*(b))$$

and

$$\mathcal{L}_{hh} = \sum_b (B_0(b)|r_0(b)|^2 + 2B_1(b)\Re(r_0(b)r_1^*(b))).$$

The values of the constant and linear coefficients for each bin, r_0 and r_1 , can be computed from the values of the ratio at the bin edges as

$$r_0(b) = \frac{r(b) + r(b+1)}{2} \quad \text{and} \quad r_1(b) = \frac{r(b+1) - r(b)}{df(b)}.$$

Crucially, this depends on the ratio $r(f)$ being relatively smooth across the frequency axis. This is expected to be approximately true for all the relevant points in the posterior, potentially not if we stray from its peak.

The formulation given here works well in the absence of spin-induced orbital precession. In its presence it breaks down, since the smooth structure of the waveforms is not preserved — the residual $r(f)$, directly computed, is not smooth enough for Equation 3.6 to be a good approximation. However, this can be counteracted by performing a mode-by-mode decomposition in the orbital frame of the binary (see Section 6.1.1.3) and then approximating those ratios as smoothly varying [130].

For an example of how the waveform ratio $r(f)$ might look, see Figure 8.14, where it is computed in a specific case for the Lunar Gravitational Wave Antenna.

To our knowledge, this framework has not been applied in the case of eccentric waveforms. Doing so in the frequency domain would likely be more efficient if a decomposition of these waveforms into harmonics is performed. The relative binning approach can also be applied to time-domain analyses [162], which might be helpful in this scenario.

3.4.1.1. Noisy summary data for injections

Relative binning allows for the cheap generation of simulated summary data under the Gaussian noise assumption [196], in the context of an injection for which the true waveform parameters are known, allowing us to set $h_0 = h_{\text{true}}$.

In that case, the data d can be decomposed as $h_0 + n$, and the summary data A_0 and A_1 can be written as

$$\begin{aligned} A_0(b) &= B_0(b) + \delta A_0(b) \\ A_1(b) &= B_1(b) + \delta A_1(b). \end{aligned}$$

While $B_0(b)$ and $B_1(b)$ are real-valued by construction, $\delta A_0(b)$ and $\delta A_1(b)$ are complex-valued random variables; their phases are uniformly distributed in $[0, 2\pi)$, while their amplitudes are distributed according to a multivariate normal distribution with mean zero and covariance

$$\langle \delta A_i^*(b) \delta A_j(b) \rangle = \begin{pmatrix} B_0(b) & B_1(b) \\ B_1(b) & B_2(b) \end{pmatrix}_{ij},$$

where we define

$$B_2(b) = \int_{f \in b} \frac{|h_0(f)|^2}{S_n(f)/T} (f - f_m(b))^2.$$

Since the noise is assumed to be uncorrelated in Fourier space, the distributions of $\delta A_0(b)$ and $\delta A_1(b)$ across different frequency bins are also independent.

3.4.2. Reduced order quadrature

Whereas relative binning seeks to closely approximate the likelihood in the region where waveforms “look similar to each other”, reduced order quadrature allows us to obtain a representation which is valid for any point in the prior space [40, 179, 126].

The equations to approximate the terms in the likelihood are somewhat similar: some weights ω and ψ are pre-computed, then the waveform is evaluated on a sparse grid and multiplied by these.

$$\Re(d|h) \approx \sum_{j=1}^{N_L} h(f_j, \theta) \omega_j$$

$$(h|h) = \sum_{i=1}^{N_Q} |h(f_i, \theta)|^2 \psi_i.$$

The weights are typically computed with a greedy algorithm, which constructs a basis that will represent all waveforms sampled from the prior to a given tolerance.

Baker et al. [15] have argued that this method, while being the best currently available option to guarantee an unbiased analysis, can only be extended so far when considering next generation detectors such as the Einstein Telescope and Cosmic Explorer. This is because, when accounting for the Earth’s rotation and frequency-dependent antenna patterns, the computational burden for creating and storing the basis becomes too large.

3.5. Marginalization

If the exact dependence of the likelihood on a given parameter is known, the integral over said parameter can sometimes be performed analytically [163, 178, 189]. This procedure is generally applied to reference phase ϕ , distance d_L and reference time t_0 [178].

With some more advanced techniques, it is possible to marginalize over all extrinsic parameters in order to minimize the number of times an expensive model for the intrinsic ones is evaluated [154, 110].

Here we will outline the analytic marginalization procedure only for the relatively simple case of reference phase, as we used this in Tissino et al. [183].

3.5.1. Reference phase

Reference phase does not have any astrophysical significance, despite being crucial for a complete description of the signal [189].

As we will outline in Section 6.1.1, the emission from a compact binary can be decomposed in harmonics, each of which has a different scaling with orbital phase. If we only consider the dominant, $\ell|m| = 22$ mode of emission, the waveform’s dependence on the coalescence

phase ϕ_c is simply $h(\phi_c, \theta) = h(\phi_c = 0, \theta)e^{2i\phi_c}$. For convenience, we shall denote the waveform at a reference phase, $h(\phi_c = 0, \theta)$, simply as $h(\theta)$.

Therefore, the marginalized likelihood reads

$$\begin{aligned} \int_0^{2\pi} \mathcal{L}(d|\phi_c, \theta) \pi(\phi_c) d\phi_c &= \mathcal{Z}_N \exp\left(-\frac{1}{2}(h(\theta)|h(\theta))\right) \int_0^{2\pi} \exp(\Re(d|h(\phi_c, \theta))) \pi(\phi_c) d\phi_c \\ &= \mathcal{Z}_N \exp\left(-\frac{1}{2}(h(\theta)|h(\theta))\right) I_0(|(d|h(\theta))|). \end{aligned} \quad (3.7)$$

The function I_0 appearing in the last expression is a modified Bessel function of the first kind, the solution to the following integral:

$$I_0(\sqrt{A^2 + B^2}) = \frac{1}{2\pi} \int_0^{2\pi} \exp(A \cos \phi + B \sin \phi) d\phi.$$

3.5.1.1. Recovering phase from marginalized samples

Since the analytical likelihood is known, it is possible to sample the phase after having performed a run with a phase-marginalized likelihood [178].

For any given sample θ from a phase-marginalized posterior we can write the posterior as a function of ϕ_c , conditioned on all other parameters θ : the integrand in Equation 3.7.

This only requires one evaluation of the intrinsic waveform, since the values are all related by a simple known phase shift; it therefore cheap to explicitly compute the inverse CDF of this distribution and extract a sample from it.

Part II.

Anatomy of a gravitational wave signal

4. Tests of General Relativity with GW250114

In this chapter, we present the results of the discovery paper for the event GW250114, the clearest to date, which enabled us to perform a high-significance test of Hawking’s area law, as well as black hole ringdown spectroscopy [2].

The current chapter and the two following it will discuss the “anatomy” of a gravitational wave signal, the way it is parameterized, and the way it appears in our detectors, using GW250114 as an example wherever possible, but also with an eye toward low frequencies, which we will then expand on in Chapter 8. The current chapter will focus on ringdown, as this is what allowed for the most prominent scientific results with GW250114. Chapter 5 will describe the intrinsic modelling of a gravitational wave signal, with parameters such as the masses and spins of the two component objects, eccentricity, tidal effects. In Chapter 6 we will bring the signal from source to detector, discussing the impact of spin-induced precession, detector motion and antenna patterns.

4.1. The signal

On January 14th, 2025, the LIGO-Virgo-KAGRA collaboration detected the clearest binary black hole merger to date: its signal to noise ratio was between 77 and 80.

The signal is clearly visible in the data with minimal processing: Figure 4.1 shows whitened strain data (see Section 3.2.3) from the Hanford and Livingston detectors.

We will attempt to give an outline of the essential physics of binary mergers, starting from what can be gleaned from the “bare” data without the need to invoke complex models [52].

We can visually identify the sections of the waveform emitted by a black hole coalescence:

- an *inspiral*, whose amplitude and frequency rise over time;
- a *merger*, where the signal amplitude peaks;
- a *ringdown*, characterized by strongly damped oscillations.

Indeed, if we overlay the general relativistic prediction for a binary merger (Figure 4.2) we can see that it matches very closely.




Figure 4.1.: Whitened strain for GW250114 in the Livingston and Hanford detectors. Adapted from Abac et al. [2].




Figure 4.2.: Upper panel: whitened strain for GW250114 in the Livingston and Hanford detectors, along with modelled and unmodelled waveform reconstructions. Lower panel: a time-frequency spectrogram of the whitened data. Adapted from Abac et al. [2].

4.2. Ringdown

After the coalescence of two Kerr black holes, the remnant quickly settles into a quiescent state, emitting gravitational radiation in a process called *ringdown*.




Figure 4.3.: Whitened strain in the ringdown of GW250114. The data from the Hanford detector was shifted in time and inverted to account for its different position and antenna patterns compared to Livingston.

The qualitative behaviour of this regime, a damped sinusoid, is apparent in Figure 4.3. Indeed, the dominant late-time behavior of this radiation can be described, at the perturbative level, in terms of quasi-normal modes (QNMs), whose complex frequencies $\omega_{\ell mn} = 2\pi f_{\ell mn} - i/\tau_{\ell mn}$ are well-understood functions of the black hole properties: mass M_f and spin magnitude χ_f .¹

For example, a remnant black hole with $M_f = 68.4M_\odot$ and $\chi_f = 0.68$ has a fundamental frequency of $f_{220} = 248$ Hz and a damping time of $\tau_{220} = 4.13$ ms.

The fact that $f_{220} \approx 1/\tau_{220}$ means that the sinusoid corresponding to this mode is approximately suppressed by a factor e every cycle. This decay is comparatively slow for a mode: most other QNMs are suppressed even faster [21, figure 2.3].

In terms of these modes, the gravitational wave polarizations from a Kerr black hole's

¹Even at the perturbative level, this description is not complete, especially at early times [65].

ringdown can be expressed as [21, equation 6.12]:²

$$h_+(t) - ih_\times(t) \approx \sum_{\ell=2}^{\infty} \sum_{m=-\ell}^{\ell} \sum_{n=0}^{\infty} A_{\ell mn} \exp(i(\omega_{\ell mn}^x(t - t_{\text{ref}}) + \phi_{\ell mn}^x))_{(-2)} Y_{\ell mn}$$

Predicting the amplitudes of these modes is in general hard, doing so in the generic-orbit case is in fact an open problem [42].

Ringdown analyses may be performed at several levels of model-dependence, ranging from a purely data-driven fit with damped sinusoids to imposing a prescription for the Kerr spectrum as well as the mode amplitudes [21, table 7]. The fits we will describe in Section 4.3 are model-dependent, imposing a Kerr spectrum but not the mode amplitudes.

In order to isolate the ringdown, one generally analyzes data starting after the coalescence, with the time domain data analysis techniques described in Section 3.2.2. Choosing a start time, however, poses a conundrum: the initial spacetime at the collision is highly dynamical and nonlinear, and it only approaches the perturbative regime in which the QNM description is formulated asymptotically. This is at the core of a fundamental tension in ringdown science: the perturbative assumption is better verified the further along in time one starts, but the signal is quickly decaying and starting too late means losing valuable data.

In figure Figure 4.4 this tension is visible for GW250114: the horizontal axis in all panels is the analysis start time, expressed in units of the characteristic timescale corresponding to the detector-frame remnant mass, $t_{M_f} = (1+z)GM_f/c^3 \approx 337\mu\text{s}$. The SNR quickly decreases with time, and the errors on the measured parameters — frequencies and damping times of the fundamental ringdown mode and its first overtone — accordingly increase with start time.

While this problem remains here as in all ringdown analyses, another fact to notice in Figure 4.4 is the very high magnitude of the ringdown-only SNR. As far as $20t_{M_f}$ after the peak it is of the order of 10, more than the entirety of some other gravitational wave signals.

4.3. Detection of the 221 overtone

The exceptional SNR of GW250114 allowed for the confident detection of the $\ell mn = 221$ overtone in the data. Its amplitude was constrained to be non-vanishing: Figure 4.5 illustrates posterior distributions on the 220 and 221 amplitudes. There is a temporal

²This expression only considers prograde modes, which co-rotate with the remnant black hole’s spin; retrograde ones are hardly excited for quasi-circular binaries.




Figure 4.4.: Frequencies, damping times and SNRs as a function of analysis start time in the ringdown of GW250114. Adapted from Abac et al. [2].

window, between 6 and 8 t_{M_f} after the peak, during which the overtone amplitude is constrained to be larger than 0 at more than 3σ .




Figure 4.5.: Posterior distributions for the amplitudes of the fundamental and first overtone of GW250114. More details in the text. Adapted from Abac et al. [2].

The grey bands in the figure indicate the theoretical expectation for the exponential decay of the two modes. They are obtained by considering the posterior samples at the reference time for that analysis ($6t_{M_f}$ for the 221, $10.5t_{M_f}$ for the 220), and evolving them backward and forward in time according to their respective amplitude and decay time. The grey band is then computed by considering the HPD intervals with 50 and 90% credibility (see Section 2.1.1.1).

4.3.1. Spectroscopic test of the Kerr metric

The confident detection of a 221 overtone enabled a spectroscopic test of General Relativity: we performed an analysis which constrained the two ringdown tones to their values with the Kerr metric, but also allowed for a deviation in the frequency and damping rate of the 221 tone:

$$f_{221} = e^{\delta f_{221}} f_{221}^{\text{Kerr}} \quad \text{and} \quad \gamma_{221} = e^{\delta \gamma_{221}} \gamma_{221}^{\text{Kerr}}.$$

This allowed the two modes to freely vary, independently of each other. In the right panel of Figure 4.6 we show the posterior distribution on the deviation parameter δf_{221} : it is constrained to be close to zero, a null result. The prior on the deviation parameter is also

shown. We can visually compare the ratio of posterior and prior, which constitutes the evidence based on this data against a deviation from General Relativity (see Equation 2.7). The precise value of this number depends on many analysis details; qualitatively, we can say we found weak evidence against a deviation from General Relativity.

The left panel offers an alternative perspective on the same result: we show the two-dimensional posterior distribution on the frequencies $(1+z)f_{220}$ and $(1+z)f_{221}$.³ The grey diagonal band shows the range of possible combinations of values under the Kerr assumption.

The two estimated frequencies change as a function of detector-frame remnant mass $(1+z)M_f$ and spin χ_f ; these are free parameters, informed only by the ringdown, hence their wide posterior distributions. If we were to use the whole signal, as opposed to only its ringdown, to inform our estimates of these parameters, we would get precise values near the black cross.




Figure 4.6.: Spectroscopic test of the 221 QNM’s frequency with GW250114. The left panel shows the posterior distribution on the 220 and 221 detector-frame frequencies, with a grey band indicating which combinations of these are allowed within General Relativity; the predicted ratio depends on the mass and spin of the remnant, hence we also mark the maximum-likelihood values for these based on the full-signal analysis. In the left panel, we show the prior and posterior distributions on the parameter δf_{221} . PYRING and RINGDOWN refer to two independent analysis pipelines we adopted. Adapted from Abac et al. [2].

³The quantities in the plot are the measured frequencies at the detector. Due to cosmological effects we expect them to be redshifted by a factor $(1+z)$ — see also Section 6.1.3.

4.4. Area law

The possibility to independently analyze the inspiral and ringdown also provided the opportunity to validate Hawking’s area law: the statement that the horizon area of black holes in the universe must increase over time.

This law is fundamentally related to the second law of thermodynamics, $dS \geq 0$, which dictates that the entropy of a closed system is non-decreasing. The entropy S of a black hole is encoded in its surface area A [90], as $S = k_B c^3 A / 4G\hbar$.

Quantum effects at the horizon (*i.e.* Hawking radiation) can be incorporated in a generalized second law [17], though they are negligible at the scales we are probing. Probing Hawking’s law means probing its assumptions:

1. General Relativity, a metric theory of driven governed by Einstein’s equations $G_{\mu\nu} = 8\pi G c^{-4} T_{\mu\nu}$, which derive from an action in the form $S = \int R d^4x$, holds.
2. The observed objects are black holes and follow the weak censorship principle, which disallows the existence of naked singularities.
3. The null-energy condition holds, meaning that for any null vector field ξ (*i.e.* satisfying $g_{\mu\nu}\xi^\mu\xi^\nu = 0$) the stress tensor satisfies $T_{\mu\nu}\xi^\mu\xi^\nu \geq 0$.

We do not have direct access to the areas of the black holes; in order to probe Hawking’s law we must estimate them based on the parameters which affect the available data. The surface area of a Kerr black hole is given in terms of its mass M and spin magnitude χ by:

$$A(M, \chi) = 8\pi \left(\frac{GM}{c^2} \right)^2 (1 + \sqrt{1 - \chi^2}) .$$

Applying this formula to our data means assuming that GW250114 was produced by a binary systems of black holes, and both the two initial components and the remnant are described by the Kerr metric. Furthermore, we need to assume that General Relativity provides a good description of binary dynamics in the weak field regime, corresponding to the early inspiral.

We do *not*, however, need to assume that our understanding of GR is correct in its strong-field regime, at merger: the exceptionally high SNR of this signal allows us to excise the merger part of our signal, and extract the component and remnant areas from independent analyses of the early inspiral and late ringdown. The fractions of data analyzed can be seen in the inset of Figure 4.7 . In order to have a statistically sound analysis of the truncated data, we need to use time-domain techniques (Section 3.2.2) both for the inspiral and for the ringdown.

In the context of a binary coalescence, Hawking’s law translates to the statement that the remnant area should be larger than the sum of the component areas:

$$A_1 + A_2 = A_i \leq A_f.$$

The inspiral and ringdown analyses do not directly give us access to the mass M of the black holes: instead, the signal directly depends on detector frame masses, in the form $(1+z)M$ — see Section 6.1.3 — and therefore we can only probe “redshifted areas” in the form $(1+z)^2 A$. Fortunately, this does not diminish the power of our analysis, as the ratio $(A_f - A_i)/A_i$ is independent of the signal’s redshift.




Figure 4.7.: Posterior distribution for the fractional area change of GW250114. We describe the meaning of the panels in the text. Adapted from Abac et al. [2].

The lower panel of Figure 4.7 illustrates the posterior distribution on this quantity, which excludes a negative area variation to high credibility. The green histogram corresponds to the aforementioned analysis performed while independently considering the pre- and post-merger sections of the signal. The upper panel showcases the robustness of this result to varying the cut time for the inspiral; the main result we quote is the one obtained when truncating the inspiral waveform $40t_M$ before the merger and starting the ringdown analysis $10.5t_{M_f}$ after the merger.

For reference, in the lower panel we also show the constraints on the change in area obtained when coherently analyzing the whole signal (“Full signal measurement”). While this is significantly more precise, General Relativity being correct is assumed to obtain this result, so it cannot function as a test of it.

Hawking’s area theorem predicts $A_f - A_i > 0$, which we show in the lower panel as a vertical dashed line. Furthermore, if we assume that energy is conserved, an upper bound for the change in area can be found by considering two maximally-spinning black holes with mass $M_1 = M_2 = M$ and $\chi = 1$ form a minimally spinning black hole with mass $2M$ and spin $\chi = 0$. Then, we will have $A_f = 4A_i$, or $(A_f - A_i)/A_i = 3$. We show the region excluded by this argument in light grey.

We further explore the robustness of this result in Figure 4.8; the rapid decay in the ringdown signal makes the choice of cut time there significantly more impactful than that in the inspiral. Despite this, the statement $A_f - A_i > 0$ is robustly verified for a wide range of choices for the end time of the inspiral analysis, $t_<$, and the start time of the ringdown analysis, $t_>$.




Figure 4.8.: Robustness of the fractional area change of GW250114 to changes in cut times. For various choices of cut time, we show the corresponding 90% interval and median for the posterior distribution on the initial and final black hole area respectively. For the final area, we show results obtained both with an analysis including both 220 and 221 overtones, as well as with an analysis including only the 220 one. Adapted from Abac et al. [2].

5. Intrinsic parameters and waveform modeling

In this chapter we continue in the dissection of the gravitational signal emitted by the coalescence of two compact objects by considering its intrinsic parameters, those relating to the source itself and not to its relation to the observer.

This gravitational signal is often qualitatively decomposed into three sections: inspiral, merger and post-merger.

When the remnant is a black hole the post-merger is an exponentially damped ringdown, as described in Section 4.2; if this is not the case, the emission spectrum can become significantly less predictable and long-lived. Though efforts have been made toward modelling it—see e.g. [32, 30, 31] for neutron star binaries—we will not discuss this in detail here.

The rest of this chapter mainly concerns the dependence of models for the inspiral on their parameters. Some of the ones we will consider are really only suitable for the early stages of inspiral, while others are accurate enough to describe the merger, and they can be smoothly connected to a model for the post-merger. When applied to black hole binaries, the result is an Inspiral-Merger-Ringdown (IMR) waveform model.

This chapter also contains a discussion of the results we presented in Tissino et al. [182], regarding the construction and application of a fast frequency-domain surrogate for neutron star binaries, as well as the ones in Gamba et al. [81], regarding a re-analysis of the GW150914 event with a waveform model accounting for both eccentricity and spin-precession.

5.1. Component masses

The two objects giving rise to the coalescence are labelled “primary” and “secondary” based on their masses, respectively m_1 and m_2 with $m_1 > m_2$. These masses are typically measured in units of the solar mass $M_\odot \approx 1.99 \times 10^{30} \text{kg}$.

Some common combinations of these include the total mass $M = m_1 + m_2$, the mass ratio $q = m_2/m_1 \leq 1$,¹ the reduced mass $\mu = m_1 m_2 / M$ and the symmetric mass ratio $\nu = \mu / M = m_1 m_2 / M^2$.

The waveform at low frequencies is most affected by their combination called *chirp mass* $\mathcal{M}$, which is defined as follows:

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

A general Post-Newtonian expression for the phase of a compact binary reads as follows [3, eq. 1.1]:

$$\Psi(f) = 2\pi f t_c - \phi_c - \frac{\pi}{4} + \frac{3}{128\nu x^5} \left(1 + \sum_{n=2}^{n_{\max}} [\varphi_n + \varphi_n^{(l)} \log x] x^n \right). \quad (5.1)$$

Here, the expansion parameter represents the velocity of the binary: $x = v/c = (\pi M f)^{1/3}$. The chirp mass appears as a prefactor in Equation 5.1: $\nu x^5 = \nu (\pi M f)^{5/3} = (\pi \mathcal{M} f)^{5/3}$. This means it affects the phase evolution of the signal at a basic level, and its measurement will be precise if we can constrain a large number of signal cycles.

On the other hand, the Post-Newtonian coefficients φ_n ² contain information about various other contributions, such as mass ratio and spin [25]. For example, the symmetric mass ratio ν enters in the expression at 1PN order, spin-orbit effects at 1.5PN, and spin-spin effects at 2PN.

5.1.1. The component masses of GW250114

This signal was quite similar to the first-ever detection, GW150914. The inferred masses of the component black holes were in fact compatible, and the distances were similar as well. However, due to the improvements in the detectors' sensitivity over the decade between them, the precision in measurement is now significantly better. Figure 5.1 shows the posterior distribution on the component masses for these two events.

¹Conventions for this parameter differ, with some authors using $q = m_1/m_2 > 1$, but within BILBY, the standard piece of software for gravitational wave inference, it is defined to be less than 1 [151].

²A terminology note: the n -th term in the sum is called the $(n/2)$ PN contribution. For example, the $n = 5$ term is called 2.5PN.




Figure 5.1.: Posterior distribution on the component masses of GW250114. Adapted from Abac et al. [2].

5.1.2. At lower frequencies

At very low frequencies we might expect to only be able to constrain the chirp mass, with a large uncertainty on the mass ratio and all other parameters appearing later in the Post-Newtonian expansion such as spin. While this is qualitatively true, an interesting observation [185] is that low frequency posterior distributions can still exhibit strong correlations between chirp mass and mass ratio.

If chirp mass were the only parameter allowed to vary, the frequency derivative of the signal $\dot{f}$, or equivalently in the stationary phase approximation (SPA, see Section 6.2) the second derivative of the frequency-domain phase Ψ'' in Equation 5.1, would be sufficient to constrain it. However, in general, fixing this frequency derivative will determine a surface — for example, a 1D curve in the $(\mathcal{M}, \nu)$ space, or a 2D surface in the $(\mathcal{M}, \nu, \chi_{\text{eff}})$ space.

Toubiana et al. [185] show this curve in the $\mathcal{M}, \nu$ plane, in the case of a stellar-mass black hole binary observed by LISA. Qualitatively, it is very “shallow”: a variation of ν from 0.2 to 0.25 (a 25% increase) only corresponds to a variation in $\mathcal{M}$ by $10^{-3} M_{\odot}$ (a 0.003% increase). Nevertheless, the scale of the variation in chirp mass is comparable to typical uncertainties in this parameter obtained with long observations, and an anticorrelation between $\mathcal{M}$ and ν is systematically observed.

5.1.3. Prior distribution

Typically, sampling is performed in terms of the variables $\mathcal{M}$ and q , but the prior distribution is required to be uniform in the (m_1, m_2) plane. The equivalent distribution can then be obtained by computing the Jacobian of the transformation [67]:

$$\pi(\mathcal{M}, q) = \mathcal{M} \frac{(1+q)^{2/5}}{q^{6/5}}.$$

When sampling, it is also important to identify which object is the primary and which is the secondary [83]. The distinction would be well-specified if we had access to the exact masses of the objects, m_1 and m_2 , but due to measurement error we do not: any given posterior sample will describe two objects with masses m_a and m_b and corresponding spins $\vec{\chi}_a$ and $\vec{\chi}_b$.

The current standard approach is to establish the correspondence $a \iff 1$ and $b \iff 2$ based on the criterion that $m_a > m_b$, on a sample-by-sample basis. If the masses are well-distinguished (*i.e.* the mass ratio is considerable) this can be reliable, but for comparable masses and low SNRs it is prone to bias inference. A generic, nonparametric and non-data-driven labelling procedure is not currently available. A practical stopgap approach is to discuss results based on relabeling invariant parameters.

5.2. Spin and orbital angular momentum

In a generic binary with orbital angular momentum $\vec{L}$, the component objects will have two spin vectors $\vec{S}_{1,2}$. These three vectors will generally not be aligned, neither to each other nor with the observation direction $\vec{N}$.³

In principle we could use Cartesian coordinates to describe these spins; however in gravitational wave analysis it is more convenient use the following decomposition:

3. $a_1 = |\vec{S}_1|/m_1^2$, the dimensionless *magnitude* of the primary object's spin;
4. $a_2 = |\vec{S}_2|/m_2^2$, the dimensionless *magnitude* of the secondary object's spin;
5. $\theta_1 = \arccos(\vec{S}_1 \cdot \vec{L})$, the tilt angle between the orbital angular momentum and the primary object's spin;
6. $\theta_2 = \arccos(\vec{S}_2 \cdot \vec{L})$, the tilt angle between the orbital angular momentum and the secondary object's spin;

³In this section we shall use this letter for the observation direction to conform with the convention, though elsewhere we use m .

7. $\phi_{12} = \arccos(\vec{S}_1^\parallel \cdot \vec{S}_2^\parallel)$, the azimuthal angle between the projections of the component spins in the orbital plane, $\vec{S}_i^\parallel = \vec{S}_i - \vec{S}_i \cdot \hat{L}$;
8. ϕ_{JL} , the azimuthal angle of $\vec{L}$, which parameterizes rotations around the $\vec{J}$ axis.

These six parameters describe the two spin vectors in terms of their relation to the orbital angular momentum. The orientation of the total angular momentum is an extrinsic parameter, since it describes the geometric relation of the source and the observer: it is described by

1. $\theta_{JN} = \arccos(\vec{J} \cdot \vec{N})$, the angle between the total angular momentum $\vec{J} = \vec{L} + \vec{S}_1 + \vec{S}_2$ and the observation direction $\vec{N}$.
2. ϕ_0 , the orbital phase, the azimuthal angle of the observation direction $\vec{N}$ around the orbital angular momentum.

Also see Romero-Shaw et al. [151] for the definitions of these parameters within BILBY.




Figure 5.2.: Diagram of the angles involved in the description of generic spins, from Du-
pletsa et al. [67].

Due to spin-induced precession, the orbital angular momentum $\vec{L}$ and the spins $\vec{S}$ are time-dependent. The total angular momentum $\vec{J}$ is not exactly conserved for an inspiralling,

spin-precessing binary [27], but we can consider a reference frame aligned to its value early in the evolution of the binary, $\vec{J} = \lim_{t \rightarrow -\infty} \vec{J}(t)$, so its angle with the observation direction θ_{JN} is a well-defined parameter [144]. The magnitudes of the spins a_i are exactly conserved. The four remaining parameters — two tilts and two azimuthal angles — measurably evolve through spin-precession, and thus need to be specified at a given reference frequency.

The angle between orbital angular momentum and observation direction, or *inclination angle*, is denoted as $\iota = \arccos(\vec{L} \cdot \vec{N})$. In the aligned-spin case it is constant and equal to θ_{JN} , but in general it is a time-dependent quantity.

The best-constrained spin-related quantity is generally the effective spin parameter

$$\chi_{\text{eff}} = \left(\frac{m_1 \vec{\chi}_1 + m_2 \vec{\chi}_2}{m_1 + m_2} \right) \cdot \hat{L}.$$

The extent to which precessing spin affects a binary is often characterized through the precessing spin parameter

$$\chi_p = \max(\chi_1^\perp, q(3q+4)/(4q+3)\chi_2^\perp),$$

where $\chi_i^\perp = |a_i \sin(\theta_i)|$ denotes the magnitude of the in-plane component of the spins.




Figure 5.3.: Effective and precessing spin parameters for GW250114, from Abac et al. [2].

In Figure 5.3 we show the posterior distribution on these parameters for GW250115, obtained with several different waveform approximants. Qualitatively, χ_{eff} was slightly negative and χ_p was small. The apparent decrease in posterior mass for $\chi_p \rightarrow 0$ is an artefact

of the prior used on the component spins (Section 5.2.1), whose distribution in terms of χ_p is shown in grey: the data prefer small values of χ_p . Note how the scale of the plot for both components could in principle be extended to a magnitude of 1.

Modelling spin precession is generally accomplished through a change in reference frame: a waveform is obtained in a co-precessing frame, whose z axis is the orbital angular momentum, and then transformed to an inertial frame by means of a time-dependent transformation, which can efficiently be computed with a Post-Newtonian approximation. We give more details on the procedure in Section 6.1.1.3.

5.2.1. Prior distribution

The prior to adopt for spin variables is a matter of convention. Within BILBY [151], it is defined so that the dimensionless spin magnitudes a_i are uniformly distributed in the range $[0, 0.99]$, while the spins are uniformly distributed on the sphere, meaning that ϕ_{12} , ϕ_{JL} , $\cos \theta_1$, $\cos \theta_2$ are uniformly distributed.

If one wishes to restrict sampling to the aligned-spin case only, care must be taken to use compatible priors: given a prior $\pi(|\chi|)$ on the spin magnitude and a prior $\pi(\cos \theta)$ on the cosine of the tilt, the equivalent prior on the aligned spin component χ_z is

$$\pi(\chi_z) = \int_0^1 da \int_{-1}^1 d \cos \theta [\pi(|\chi|) \pi(\cos \theta) \delta(\chi_z - |\chi| \cos \theta)] .$$

If the prior on the spin magnitude $\pi(|\chi|)$ is taken to be uniform in the range $[0, \chi_{\max}]$, then the result of this integral is [110, eq. A7]:

$$p(\chi_z) = -\ln(|\chi_z|/\chi^{\max})/2\chi^{\max} . \quad (5.2)$$

5.3. Eccentricity

In Newtonian gravity, the two-body problem is analytically solvable. The corresponding trajectories are all planar, and they can be expressed in polar coordinates as [103]:

$$r = \frac{p}{1 + e \cos f} ,$$

where f is a polar angle in the orbital plane, measured from the center of mass of the system, while $r = |r_1 - r_2|$ is the distance between the two objects. The motion of either body can easily be recovered from this equation: their position vectors measured from the center of

mass will always have opposite signs, and magnitudes $r_1 = qr/(1+q)$ and $r_2 = r/(1+q)$, where q is the mass ratio. The constants of motion p and e are called *semilatus rectum* and *eccentricity*. The former has the dimensions of a length and it characterizes the orbit's scale; the latter is a pure number, which determines the type of conic section the orbit describes:

- $e = 0$ is a circular orbit;
- $0 < e < 1$ is an elliptical orbit;
- $e = 1$ is a parabolic trajectory;
- $e > 1$ is a hyperbolic trajectory.

Systems with $e < 1$ are *bound*, they repeat their motion and have negative total energy. Let us define some parameters for the elliptical case: the points (x, y) of an ellipse with semimajor axis a and semiminor axis b obey

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$$

In terms of these variables, the *eccentricity* e is given by $e^2 = 1 - (b/a)^2$, the *focal distance* c by $c = \sqrt{a^2 - b^2}$, and the *semilatus rectum* p by $p = a\sqrt{1 - e^2}$.

We have previously used the *true anomaly* f as a polar angle: it is the angle from the periastron to the orbiting body, measured using the *focus of the ellipse* as the center. There are some other choices: the *eccentric anomaly* E is the angle from the periastron to the *projection* of the orbiting body onto the smallest circle containing the ellipse, measured using the *center of the ellipse* as the center. The *mean anomaly* M is an angle increasing proportionally to time from a periastron pass to the next.

i Note 5: Converting between eccentric parameters

Mean anomaly and eccentric anomaly are related through Kepler's equation, which is transcendental and therefore must be solved numerically if we want to compute E as a function of M :

$$M = E - e \sin E$$

True anomaly and eccentric anomaly are related through the equation:

$$(1 - e) \tan^2(f/2) = (1 + e) \tan^2(E/2)$$

but solving it is numerically unstable (the tangent diverges), so there is a better way [33]: define $\beta = e/(1 + \sqrt{1 - e^2})$, then

$$f = E + 2 \operatorname{atan2}(\beta \sin E, 1 - \beta \cos E)$$

where $\text{atan2}(y, x) = \arctan(y/x)$ is the two-argument tangent function, which is also stable for $x = 0$.

Finally, one may close the circle of conversions with the following relation:

$$M = \text{atan2}(-\sqrt{1 - e^2} \sin f, -e - \cos f) + \pi - e \frac{\sqrt{1 - e^2} \sin f}{1 + e \cos f}$$

In Figure 5.4 we illustrate how these anomalies can vary as a function of time. Also see [this animated visualization](#).




Figure 5.4.: The three anomalies for a generic orbit

They provide a useful starting point toward an understanding general-relativistic eccentric binaries, but their characteristics do not fully generalize: when gravity is strong eccentricity is not a constant of motion, nor is it covariantly defined; orbits are also not exactly elliptical, due to both perihelion precession and gravitational wave emission [75].

Furthermore, the strict classification of Newtonian binaries breaks down in General Relativ-

ity: eccentricity decreases in a process known as *circularization*, and an unbound system can become bound if the gravitational wave emission during a hyperbolic encounter is enough to make its total energy negative.⁴

It is possible to define a notion of eccentricity based on the evolution of the emission frequency of the binary. This frequency is monotonic when the orbit is quasi-circular, while for eccentric orbits it exhibits peaks and troughs on the orbital timescale [161].

Within the TEOBResumS-DAlf waveform model [26, 7], eccentricity and relativistic anomaly are used as an initial condition for the EOB variables, which are then evolved with fully relativistic prescriptions (see also Section 5.5.2).

Gamba et al. [81] performed a first-of-its-kind reanalysis of the event GW150914 with a model including both eccentricity and precessing spins. Beyond the regular 15 parameters describing a spin-precessing black hole binary,⁵ this analysis required two more: eccentricity and (true) anomaly. The constraints on these two parameters are shown in Figure 5.5.

5.3.1. At lower frequencies

Most of the gravitational waves detected so far, in the audio band, have had no measurable eccentricity. GW250114 was one such example: parameter estimation performed with an eccentric waveform [26] model prefers eccentricity values very close to zero, as measured at a reference frequency of 13.33Hz.⁶

Figure 5.6 shows the posterior: the eccentricity is bounded to be lower than roughly 0.02.

Despite this being the best constraint on the eccentricity of a black hole binary merger to date, based on this observation alone we cannot exclude the possibility that this binary had significant eccentricity at lower frequencies. Figure 5.7 shows the posterior distribution for GW250114 evolved backward in time. This is done with a simple Post-Newtonian approach, namely Peters’ equations [138, 119], as it only serves to illustrate the point. It would be possible to also account for the impact of spin-precession, as described by Fumagalli et al. [78].

⁴In GR this Newtonian condition should be phrased as: “make the total energy of the system E smaller than the sum of the rest masses of the components, M .”

⁵For convenience we give the list here: 2 masses, 6 independent spin components, 7 extrinsic parameters (inclination θ_{JN} , phase ϕ_0 , 2 sky coordinates, luminosity distance d_L , polarization ψ , time of coalescence t_0 .)

⁶This specific value was chosen to correspond to the start of the generated waveform in the time domain. The frequency band used for the analysis started from 20Hz, but the waveform needed to be started earlier in order to account for the contributions at 20Hz of the $\ell = 3$ modes, which appear at a frequency $3/2$ times higher.




Figure 5.5.: Eccentricity and true anomaly posterior for GW150914, from Gamba et al. [81].




Figure 5.6.: Eccentricity posterior distribution for GW250114. A vertical line shows the location of the 90% upper limit on eccentricity.




Figure 5.7.: Backward-extrapolated eccentricity evolution for GW250114.

While the limits and analysis quoted in the discovery paper [2] used a uniform prior in $e_{13.33\text{Hz}}$, a conservative choice, we also show results obtained with a log-uniform prior $p(e_{13.33\text{Hz}}) \propto 1/e_{13.33\text{Hz}}$. The qualitative scenario is similar in both cases: eccentricities on the order of $e \sim 0.5$ are plausible in the deci-hertz band. However, the prior significantly affects whether these lie on the edge of the distribution or along its bulk.

This is not surprising: the data is not very informative on the $e_{13.33\text{Hz}} \sim 0$ tail of the distribution, so we get back out of the analysis the assumptions we started with. This result can lead us to reflect on the meaning of the prior on eccentricity: a flat distribution in $e_{13.33\text{Hz}}$ is not a neutral statement, but instead one which heavily favors eccentric binaries at lower frequencies. For an in-depth discussion of this effect, as well as an exploration of how one might use astrophysically-motivated priors instead, see Clarke et al. [49].

5.4. Tidal parameters

Objects such as neutron stars can be deformed by the tidal gravitational field caused by the presence of a nearby object, such as the binary companion. This deformation enhances the gravitational attraction between the two objects, but it is only effective at relatively short separations: the net effect, therefore, is to leave the early inspiral almost unchanged, while accelerating its late stages.

In the static case, this tidal field can be expressed as $\mathcal{E}_{ij} = C_{0i0j}$, where C is the Weyl curvature tensor [124, eq. 13.50], the trace-free component of the Riemann curvature tensor. At leading order, this tidal field will induce a quadrupole moment Q_{ij} , given by

$$Q_{ij} = -\Lambda M^5 \mathcal{E}_{ij}$$

where Λ is the *tidal deformability* or *polarizability* parameter, written as

$$\Lambda = \frac{2}{3} \kappa_2 \left(\frac{Rc^2}{GM} \right)^5 \quad (5.3)$$

where R is the radius of the object, while κ_2 is its $\ell = 2$ dimensionless tidal Love number, which for neutron stars is generally on the order of 0.1.

The main tidal contribution to the waveform is through the effective tidal parameter, which enters the expansion at 5PN order [190]:

$$\tilde{\Lambda} = \frac{8}{13} \left[(1 + 7\nu - 31\nu^2)(\Lambda_1 + \Lambda_2) + \sqrt{1 - 4\nu}(1 + 9\nu - 11\nu^2)(\Lambda_1 - \Lambda_2) \right]$$

while secondary contributions arise depending on the tidal deviation parameter

$$\delta\tilde{\Lambda} = \frac{1}{2} \left[\sqrt{1-4\nu} \left(1 - \frac{13272}{1319}\nu + \frac{8944}{1319}\nu^2 \right) (\Lambda_1 + \Lambda_2) + \left(1 - \frac{15910}{1319}\nu + \frac{3380}{1319}\nu^2 \right) (\Lambda_1 - \Lambda_2) \right].$$

While their expressions are intimidating, note that several terms vanish in the equal-mass case, where $\nu = 1/4$; if that is true then $\tilde{\Lambda} = f(\nu)(\Lambda_1 + \Lambda_2)$ and $\delta\tilde{\Lambda} = g(\nu)(\Lambda_1 - \Lambda_2)$ for some functions f and g . Furthermore, if the masses are equal and so are the tidal parameters ($\Lambda_1 = \Lambda_2 = \Lambda$), then we have $\Lambda = \tilde{\Lambda}$ and $\delta\tilde{\Lambda} = 0$.

The analysis so far has focused on the adiabatic, quadrupolar case: assuming that only the dominant oscillation mode of a neutron star is excited, and that its time evolution follows the forcing term (*i.e.* the other star's tidal field).

In reality, the tidal deformation of a neutron star is more accurately described through a decomposition including the time-dependent excitation of different modes. There are several modes of oscillation one might consider, which have been deeply discussed in the field of astroseismology, [57, 171, 200], and a full discussion of them is beyond the scope of this work. The amplitude Q_n of each of these modes, which we index with an integer n , will oscillate following a driven harmonic oscillator [109] in the form

$$\ddot{Q}_n + \frac{\dot{Q}_n}{t_{\text{diss}}} + \omega_n^2 Q_n^2 = f_n(Q_n, t, \text{orbit}),$$

for some values of the dissipative timescale t_{diss} and angular frequency ω_n . The forcing term f_n will generally depend on the relative position of the companion star. In the quasi-circular case this can be approximated as a sinusoid whose period scales inversely with the angular frequency of the binary ω_{orb} , but for eccentric binaries high-frequency modes may be excited even though the orbit-averaged angular frequency is comparatively low.

The adiabatic assumption is equivalent to the requirement that $1/t_{\text{diss}} \ll \bar{\omega}_{\text{orb}}$, as well as $\omega_n \gg \bar{\omega}_{\text{orb}}$ [77]. The first constraint is generally expected to be well-satisfied, both when discussing losses due to gravitational wave emission [172] as well as due to friction [149].

Relaxing the $\omega_n \gg \bar{\omega}_{\text{orb}}$ assumption means considering the effect of *dynamical* (also known as *resonant*) tides. These can impact the phasing at low frequencies, especially in the case of eccentric binaries [79]. Considering oscillations with angular structure at higher orders than the quadrupole is also possible [101]. In any case, none of these higher-order corrections are required to explain current observations of neutron star mergers: the clearest to date, GW170817, only provided comparatively weak constraints on the leading-order tidal parameters Λ [51].

Love numbers are generally expected to vanish for black holes, and to be nonzero for other compact objects such as neutron stars. GW250114 provided an opportunity to test this prediction: if it had a measurable deviation from $\tilde{\Lambda} = 0$, it would have indicated a potential violation of General Relativity or some other missing aspect in our description of this binary, such as the components not being black holes but other compact objects. Andrés-Carcasona and Santoro [8] have shown this not to be the case within current experimental uncertainties, constraining $\tilde{\Lambda}$ to be less than 155 (with uniform priors) or 35 (with log-uniform priors) for this event.

5.5. Waveform modelling

General relativity is a complex non-linear theory, and in order to analyze signals from binaries with strong gravity we must understand its predictions precisely, and be able to compute them quickly.

There are several approaches to modelling the waveform emitted by a compact binary inspiral, with varying degrees of computational complexity, completeness in the inclusion of physical effects, and accuracy. For example, while all the approaches discussed here can in principle include any amount of modes in the spherical harmonic decomposition of binary emission (see Section 6.1.1), some specific models only include the dominant one.

Often, these approaches are calibrated or validated against full Numerical Relativity simulations. These have been possible for around 20 years as of this writing [145, 86], and despite being a gold standard suffer from their own set of uncertainties.

Gravitational waves can be extracted from them based on the Weyl scalar $\Psi_4 = C_{abcd}m^am^bn^cn^d$, where m and n are component vectors of a null tetrad while C is the Weyl tensor. At large separations from the binary, this is related to the gravitational time-domain polarizations $h_+ - ih_\times$ by

$$-\Psi_4 = \ddot{h}_+ - i\ddot{h}_\times,$$

so this scalar must be computed at the edges of the simulation, integrated twice in the time variable, and then decomposed into spherical harmonics [116].

We will focus on the Post-Newtonian (Section 5.5.1) and Effective One Body (Section 5.5.2) approaches to waveform generation; we also note the existence of Numerical Relativity surrogates (*e.g.* NRSur7dq4, presented by Varma et al. [187]) and phenomenological models

(*e.g.* IMRPhenomXHM, presented by García-Quirós et al. [82]) for comparable mass binaries, and the gravitational self-force formalism for extreme mass ratios (*e.g.* the second-order self-force calculations presented by Pound et al. [143] and implemented in the Fast EMRI Waveforms software package [104]).

5.5.1. Post-Newtonian

Post-Newtonian waveforms are obtained by analytically expanding Einstein’s equations in powers of v/c . Thanks to the stationary phase approximation, this can lead to an analytical expression for the amplitude and phase as a function of frequency.

These models are generally only accurate for a description of the early inspiral.

Buonanno et al. [36] give a comparison of some commonly used Post-Newtonian waveform approximants, including the frequency-domain one TaylorF2, whose frequency-domain waveforms read [123]:

$$\tilde{h}(f) = \sqrt{\frac{5}{24}} \frac{c}{\pi^{2/3} d_L} \left(\frac{GM}{c^3} \right)^{5/6} f^{-7/6} \left(1 + \sum_{n=2}^{n_{\max}} A_n (\pi M f)^n \right) e^{i\Psi(f)},$$

where the phase Ψ is in the form reported in Equation 5.1.

Post-Newtonian approximants including eccentricity also exist, though their validity is limited to the small-eccentricity regime.

5.5.2. Effective One Body

Effective One Body waveforms are obtained by mapping the relativistic two body problem to the motion of a test particle in a deformed metric. The following description will mainly describe the components of the TEOBResumS family of models, though similar considerations also apply to other models, such as those in the SEOB family.

The problem is described in a Hamiltonian setting, in terms of the dimensionless variables $(r, p_r, \varphi, p_\varphi)$, which are rescaled relative polar coordinates in the orbital plane.

The effective Hamiltonian of the problem is given by [129]:

$$\hat{H} = \frac{H}{\mu} = \frac{1}{\nu} \sqrt{1 + 2\nu(\hat{H}_{\text{eff}} - 1)}$$

where the effective Hamiltonian contains a plethora of physical effects and Post-Newtonian contributions. For the TEOBResumS model, it is given by

$$\hat{H}_{\text{eff}} = \sqrt{p_{r_*}^2 + A \left(1 + \frac{p_\varphi^2}{r_c^2} + 2\nu(4 - 3\nu) \frac{p_{r_*}^4}{r_c^2} \right)} + p_\varphi (G_S \hat{S} + G_{S_*} \hat{S}_*),$$

where the following quantities appear:

- the potential A is a generalization of the $1 - 2GM/rc^2$ term in the Schwarzschild metric, which can also include corrections due to e.g. tidal effects; though it does not explicitly feature here, the Hamiltonian also depends on the B potential, which at leading order reads $1 + 2GM/rc^2$;
- the radial momentum p_{r_*} expressed with respect to the tortoise coordinate $r_* = \int \sqrt{B/A} dr$;
- the centrifugal radius r_c , introduced in Damour and Nagar [59];
- a spin term, which depends on the variables $\hat{S} = (S_1 + S_2)/M^2$ and $\hat{S}_* = (qS_1 + S_2/q)/M^2$;

The potentials A and B are computed with Padé resummation in order to improve stability where the Taylor approximant would diverge.

Since we are working within General Relativity the equations of motion are not conservative: they include *radiation reaction* terms which affect the evolution of the momenta. These terms also need to be computed to high order and Padé-resummed.

i Note 6: Padé approximation

A (n, m) Padé approximant [134] to a function $f(x)$ is a ratio of polynomials, in the form:

$$P_m^n(x) = \frac{\sum_{i=0}^n a_i x^i}{1 + \sum_{i=1}^m b_i x^i},$$

where the coefficients are determined such that they agree with the Taylor series:

$$\frac{d^r}{dx^r} f(0) = \frac{d^r}{dx^r} P_m^n(0)$$

for all $r \leq n + m$. In fact, a Taylor series is a special case of a Padé approximant with $m = 0$.

One advantage of using Padé approximants is that, for large values of x , they have a lower asymptotic order, and thus tend towards less “extreme” behavior, as Figure 5.8 shows.




Figure 5.8.: A Padé approximant for the sine function. Even when the approximation breaks down, the Padé approximant outputs values somewhat close to those of the actual function, while the Taylor series quickly diverges.

The waveform obtained by computing a solution to the EOB dynamics is then completed with corrections for the “plunge” regime [60], and — for black hole binaries — a smooth connection to the ringdown [58], which is modelled as a superposition of quasi-normal modes as described in Section 4.2. The treatment for the post-merger regime of other compact objects such as neutron stars is significantly more complex, as the phenomenology is much more variable and not currently constrained by observational data. Nevertheless, approaches have been proposed to handle it [30].

A numerical solution to the EOB evolution equations is not always necessary: Nagar and Retegno [127] showed that a post-adiabatic approximation allows for the analytic solution to these equations. Gamba, Bernuzzi, and Nagar [80] further applied a stationary phase approximation (SPA, see Section 6.2) to these waveforms, allowing for direct generation in the frequency domain.

The SPA approach is not always feasible: for example, eccentric-hyperbolic binaries do not satisfy the requirements for the SPA, and their waveforms need to be computed in the time domain [128, 81].

The EOB approach is also amenable to being extended with beyond-GR parameters, see e.g. Chiamello et al. [47].

5.5.3. Surrogate models: the `mlgw_bns` example

The evaluation of gravitational waveform models is commonly a bottleneck for the likelihood, therefore significant efforts have been devoted to the development of *surrogate models*, which reproduce the results from a computationally expensive model.

As an example, we show results from `mlgw_bns` [182], a surrogate of the frequency-domain inspiral-only model `TEOBResumSPA` [80] for aligned-spin neutron star binaries with tidal effects — a 5-dimensional parameter space, with parameters $(q, \chi_1, \chi_2, \Lambda_1, \Lambda_2)$, since all other parameters can be modeled analytically.

As a surrogate, it is trained on the base model, which can be evaluated at arbitrary points in the parameter space. The surrogate’s accuracy depends on the number of training data points, and it is measured in terms of the mismatch

$$\bar{\mathcal{F}} = 1 - \max_{t_0, \phi_0} \frac{(h_{\text{true}} | h_{\text{surr}}(t_0, \phi_0))}{\sqrt{(h_{\text{true}} | h_{\text{true}})(h_{\text{surr}} | h_{\text{surr}})}} \quad (5.4)$$

where the scalar products are computed as in equation Equation 3.3, and generally maximized over reference time and coalescence phase — this is equivalent to allowing the surrogate model to have a different parameterization for timing and phase than its baseline. This is generally acceptable, since timing and phase have no astrophysical meaning.




Figure 5.9.: Reconstruction accuracy for `mlgw_bns` as a function of the number of training waveforms N .

Whereas the mismatch is a global metric, reconstruction accuracy can also be locally measured by directly comparing the predictions with those of the base model. We show this in

Figure 5.10: the high accuracy achieved at low frequencies is by construction, as the model is constructed by incorporating an analytic Post-Newtonian baseline and only fitting the residual error.




Figure 5.10.: Reconstruction accuracy for `mlgw_bns` in the frequency domain.

The surrogate model should also be faster to evaluate: this is indeed the case, as Figure 5.11 shows. The number of grid points in frequency has a significant impact on the evaluation time for all models, as is inevitable given the increased number of floating point operations. Therefore, we separate out the constant-time contribution to the evaluation time from the one scaling with the number of grid points.




Figure 5.11.: Evaluation time for `mlgw_bns` compared to its reference model `TEOBResumSPA` and some others. For every model, we also show a fit of the waveform evaluation time t with a linear function of the number of frequency grid points N , so that $t = a + bN$.

6. Extrinsic parameters and signal projection

In this chapter, we outline the necessary ingredients for a full description of the response of a gravitational wave detector to a compact binary coalescence signal, including the parameters describing the geometry of their relative position (Section 6.1), the way the signal can be quickly computed in the frequency domain (Section 6.2), the antenna patterns of ground- and Moon-based detectors (Section 6.4), and the calibration of ground-based detectors (Section 6.3).

6.1. Extrinsic parameters

6.1.1. Phase and inclination angle

As we introduced in Section 5.2, the reference orbital phase of the signal ϕ_0 and the inclination angle ι define the position of the observer relative to the orbital plane. In order to better understand their effect, the rest of this section will describe them in relation to the way a waveform is decomposed in spherical harmonics.

6.1.1.1. Spherical harmonic decomposition

The two time-domain polarizations of a gravitational wave, h_+ and $h_\times$, can be decomposed into spin-weighted spherical harmonics as

$$h_+(t) - ih_\times(t) = \frac{1}{d_L} \sum_{l=2}^{\infty} \sum_{m=-l}^l h_{\ell m}(t; \theta_{\text{int}}) {}^{(-2)}Y_{\ell m}(\iota, \phi_0). \quad (6.1)$$

This decomposition is performed around the *orbital* angular momentum, which is not conserved in the generic, spin-precessing case. If the binary is spin-aligned, on the other hand, it is also symmetric about the orbital plane, which leads to the relation $h_{\ell m} = (-1)^\ell h_{\ell -m}^*$ [144]. For the expressions of the spin-weighted harmonics, see Note 7.

To a good approximation, the time-domain phase of each harmonic component satisfies

$$\phi_{\ell m}(t) \approx \frac{m}{2} \phi_{22}(t),$$

which can also be written in the frequency domain as [82]

$$\phi_{\ell m}(f) \approx \frac{m}{2} \phi_{22} \left(\frac{2f}{m} \right) .$$

i Note 7: Spin-weighted spherical harmonics

The spin-weighted spherical harmonics are functions of the orientation of the source, parametrized as ι (inclination angle, between the observation direction and the source angular momentum) and ϕ_0 (initial phase of the source's rotation). Explicitly, they are [5]:

$$({}^{-s})Y_{\ell m} = (-1)^s \sqrt{\frac{2\ell+1}{4\pi}} d_{m,s}^{\ell}(\iota) e^{im\phi_0}$$

where $d_{m,s}^{\ell}(\iota)$ is called a Wigner d -function, given by

$$d_{m,s}^{\ell}(\iota) = \sum_{k=k_1}^{k_2} \left[\frac{(-1)^k \sqrt{(\ell+m)!(\ell-m)!(\ell+s)!(\ell-s)!}}{(\ell+m-k)!(\ell-s-k)!k!(k+s-m)!} \right. \\ \left. (\cos(\iota/2))^{2\ell+m-s-2k} (\sin(\iota/2))^{2k+s-m} \right], \quad (6.2)$$

where $k_1 = \max(0, m-s)$ and $k_2 = \min(\ell+m, \ell-s)$.

Note that in the gravitational wave case the label for the harmonics Y is -2 , but the parameter s in the d -function is equal to $+2$.

i Note 8: Identities

These harmonics form an orthogonal basis:

$$\int d\Omega ({}^{-2})Y_{\ell m} ({}^{-2})Y_{\ell' m'}^* = \delta_{\ell\ell'} \delta_{mm'} ,$$

and they satisfy

$${}^s Y_{\ell m} = (-1)^{s+m} Y_{\ell-m}^* .$$

i Note 9: The time-domain 22 wave

We derive the time-domain expression for the $\ell = 2$, $|m| = 2$ harmonic, which is the most commonly used approximation for a gravitational wave.

These are the relevant harmonics:

$${}^{(-2)}Y_{2\pm 2} = \sqrt{\frac{5}{64\pi}}(1 \pm \cos \iota)^2 e^{\pm 2i\phi}.$$

According to the expression we gave earlier, we will have

$$\begin{aligned} h(t) &= h_+(t) - ih_\times(t) = \frac{1}{d_L} \sum_{\ell=2} \sum_{|m|=2} H_{\ell m}(t) {}^{(-2)}Y_{\ell m} \\ &= \frac{1}{d_L} (H_{22}(t) {}^{(-2)}Y_{22} + H_{2-2}(t) {}^{(-2)}Y_{2-2}) \\ &= \frac{1}{d_L} \sqrt{\frac{5}{64\pi}} (H_{22}(t)(1 + \cos \iota)^2 e^{2i\phi_0} + H_{2-2}(t)(1 - \cos \iota)^2 e^{-2i\phi_0}) \end{aligned}$$

Now, we can make use of the fact that, thanks to symmetry under reflection across the orbital plane, we have $H_{\ell m} = (-1)^\ell H_{\ell -m}^*$, which in this case reduces to $H_{22} = H_{2-2}^*$ (see section II.D in Ossokine et al. [133]). Therefore, if we define $\tilde{H}_{22} = H_{22}e^{2i\phi_0}$, we will have

$$\begin{aligned} h(t) &= \frac{1}{d_L} \sqrt{\frac{5}{64\pi}} (\tilde{H}_{22}(t)(1 + \cos \iota)^2 + \tilde{H}_{22}^*(t)(1 - \cos \iota)^2) \\ &= \frac{1}{d_L} \sqrt{\frac{5}{64\pi}} (2\Re \tilde{H}_{22}(t)(1 + \cos^2 \iota) + 4i\Im \tilde{H}_{22}(t) \cos \iota) \\ &= h_+ - ih_\times, \end{aligned}$$

where $\Re$ and $\Im$ denote the real and imaginary parts.

At this point, we can identify the real and imaginary components, as well as expressing $\tilde{H}_{22} = H_{22}e^{2i\phi_0} = A_{22}(t)e^{i\phi_{22}(t)+2i\phi_0}$:

$$\begin{aligned} h_+ &= \frac{4}{d_L} \sqrt{\frac{5}{64\pi}} A_{22}(t) \frac{1 + \cos^2 \iota}{2} \cos(\phi_{22}(t) + 2\phi_0) \\ h_\times &= -\frac{4}{d_L} \sqrt{\frac{5}{64\pi}} A_{22}(t) \cos \iota \sin(\phi_{22}(t) + 2\phi_0) \end{aligned}$$

This is the same expression we get in the quadrupole, Newtonian approximation — see, for example, equations 4.3 in Maggiore [119], as long as we reabsorb the coefficients into the amplitude.

6.1.1.2. Inferring phase

The reference phase from a compact binary merger is a crucial parameter for its description, which however brings little physical insight [189]. This has led to efforts to simplify its treatment. If the signal only includes the $\ell|m| = 22$ harmonic, it is possible to derive a formulation of the likelihood which is analytically marginalized over phase, see Section 3.5.1.

Phase is also difficult to infer when using neural posterior estimation (Section 2.3.2): the network tends to be unable to infer it, and returns the prior instead, despite correctly constraining the remaining parameters. To this end, Dax et al. [63] introduced a resampling method, in which the phase is synthetically sampled by evaluating the likelihood on a grid in phase.

6.1.1.3. Twisting: the construction of spin-precessing frequency-domain waveforms

The construction of spin-precessing gravitational waveforms is generally performed by computing them in a non-inertial, precessing frame aligned with the orbital angular momentum $\vec{L}$, and then converting to the approximately inertial frame aligned with the total angular momentum $\vec{J}$ by means of a time-dependent rotation [144, eq. 3.3]:

$$h_{\ell m}^J = \sum_{m'=-\ell}^{\ell} \mathcal{D}_{m'm}^{\ell*}(\alpha, \beta, \gamma) h_{\ell m'}^L.$$

This choice is not unique, and the $\vec{J}$ vector is not exactly conserved for an inspiralling, spin-precessing system [27]; we can use its value at a given moment (such as $t \rightarrow -\infty$) as a reference.

This is useful because the gravitational waves in the precessing frame are quite close to those which would be emitted by a non-precessing binary.

The $\mathcal{D}_{m'm}^{\ell}$ are the Wigner D-matrices

$$\mathcal{D}_{m'm}^{\ell}(\alpha, \beta, \gamma) = e^{im'\alpha} e^{im\gamma} d_{m'm}^{\ell}(\beta),$$

where the $d_{m'm}^{\ell}$ are the Wigner d -function, defined in Equation 6.2. The three Euler angles α , β and γ are time-dependent. They allow for more freedom than necessary: the rotation mapping $\vec{L}$ to $\vec{J}$ is not unique. Once the angles α and β are determined to track the correct rotation, γ can be gauge-fixed such that

$$\dot{\gamma} = -\dot{\alpha} \cos \beta,$$

which enforces a lack of rotation in the precessing frame about $\vec{L}$ [27].

After some computation, one gets the following expression for the positive-frequency waveforms in the inertial frame:

$$\begin{aligned}\tilde{h}_+^J(f) &= \frac{1}{2} \sum_{\ell \geq 2} \sum_{m'=0}^{\ell} \tilde{h}_{\ell-m'}^L(f) e^{im'\gamma} \sum_{m=-\ell}^{\ell} [A_{m-m'}^{\ell} + (-1)^{\ell} A_{mm'}^{\ell*}] \\ \tilde{h}_{\times}^J(f) &= \frac{i}{2} \sum_{\ell \geq 2} \sum_{m'=0}^{\ell} \tilde{h}_{\ell-m'}^L(f) e^{im'\gamma} \sum_{m=-\ell}^{\ell} [A_{m-m'}^{\ell} - (-1)^{\ell} A_{mm'}^{\ell*}] ,\end{aligned}\tag{6.3}$$

written in terms of the mode-by-mode transfer functions

$$A_{mm'}^{\ell} = e^{-im\alpha} d_{mm'}^{\ell}(\beta)^{(-2)} Y_{\ell m} .$$

At a first approximation, the L -frame modes can be approximated with a non-spin-precessing model, but this is not completely accurate.

The angles defining the transformation need to be computed in the frequency domain as well, at a frequency $2\pi f/m'$. Post-Newtonian expressions for them can be used, such as those given in appendix G.1 of Pratten et al. [144].

i Note 10: Fourier transforming a complex signal

Starting from equation Equation 6.1, we want to obtain frequency-domain expressions for both the plus and cross polarizations, which are real-valued. They can be extracted by taking the real and imaginary part, and then Fourier transforming:

$$\begin{aligned}h_+(f) &= \frac{1}{2} [\tilde{h}(f) + \tilde{h}^*(-f)] \\ h_{\times}(f) &= \frac{i}{2} [\tilde{h}(f) - \tilde{h}^*(-f)] .\end{aligned}$$

We can recognize this structure in Equation 6.3.

6.1.2. Sky position and polarization

Within General Relativity, the strain from a gravitational wave moving along a fixed direction only has two independent polarizations. Alternative theories of gravity allow for the excitation of up to six [131], but we will focus on the GR case, where the strain tensor is decomposed as

$$h_{ij}(t) = h_+(t) e_{ij}^+ + h_{\times}(t) e_{ij}^{\times}\tag{6.4}$$

as a function of the two orthogonal unit basis tensors

$$e_{ij}^+ = u_i u_j - v_i v_j \quad \text{and} \quad e_{ij}^{\times} = u_i v_j + v_i u_j .$$

Three angles in total are required to define these two vectors. These are the sky position of the source, m , which is orthogonal to both u and v , and a *polarization angle* ψ which parameterizes rotations along the m axis.

The explicit expressions for these three vectors are as follows [131]:

$$m = \begin{pmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \end{pmatrix}$$

$$u = \begin{pmatrix} -\sin \phi \\ \cos \phi \\ 0 \end{pmatrix} \cos \psi - \begin{pmatrix} \cos \theta \cos \phi \\ \cos \theta \sin \phi \\ -\sin \theta \end{pmatrix} \sin \psi$$

$$v = \begin{pmatrix} -\sin \phi \\ \cos \phi \\ 0 \end{pmatrix} \sin \psi + \begin{pmatrix} \cos \theta \cos \phi \\ \cos \theta \sin \phi \\ -\sin \theta \end{pmatrix} \cos \psi,$$

where the sky position is parameterized by an azimuth $\phi \in [0, 2\pi)$ and a colatitude $\theta \in (0, \pi)$, which is equal to $\pi/2$ on the equator. Figure 6.1 illustrates the meaning of right ascension $ra = \phi$ and declination $dec = \pi/2 - \theta$, a standard parameterization for positions in the sky. Declination is analogous to latitude, and it is equal to 0 on the equator.




Figure 6.1.: Diagram of the right ascension and declination coordinates, from Dupletsa et al. [67].

When constraining the sky position of a given gravitational wave source, we typically do not actually care about the error in the coordinates themselves but about the sky

localization *area*, the minimal solid angle that can be chosen such that the probability that the signal lies within it is X , where X is some conventionally chosen number such as 90%; see Section 2.4.1.2 and Section 2.1.1.1 for more details.

6.1.2.1. Priors

Right ascension and declination are generally taken to have agnostic, uniform prior distributions in the sky, as gravitational waves originate from isotropically distributed extragalactic sources. In order to achieve this, the priors on $\text{dec} \in [-\pi/2, \pi/2]$ and $\text{ra} \in [0, 2\pi]$ are [151, 67]:

$$\pi(\text{dec}) = \frac{1}{2} \cos(\text{dec}) \quad \text{and} \quad \pi(\text{ra}) = \frac{1}{2\pi}.$$

While the polarization angle ψ could in principle take any value in $[0, 2\pi]$, this is redundant: as gravitational radiation has spin 2, its period for rotations around the propagation axis is π . Therefore, it is customary to only consider $\psi \in [0, \pi]$ with $\pi(\psi) = 1/\pi$.

6.1.3. Distance and redshift

In flat spacetime, gravitational wave amplitudes decay proportionally to $1/d$, where d is the distance to the source.

Like electromagnetic signals, however, binaries at cosmological distances are expected to be affected by cosmological redshift. This affects their propagation, modifying their frequencies. We generally work under the assumption that, on large enough scales, the Universe is described by a Friedmann-Lemaître-Robertson-Walker (FLRW) cosmological model, whose metric in spherical coordinates is

$$ds^2 = -c^2 dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \right].$$

The curvature parameter k can take on the values $k = -1$ (open Universe), $k = 0$ (flat universe), $k = +1$ (closed Universe). The parameter $a(t)$ is the *scale factor*.

The following prescriptions are required to consider cosmological effects in the analysis of a signal [119, section 4.1.4].

The Euclidean notion of distance d should be replaced by *luminosity distance* d_L . This is defined by the requirement that radiation flux from any source, as measured in the observer's frame, should decay proportionally to $1/d_L^2$; in FLRW cosmology it can be computed as

$$d_L(z) = c(1+z) \int_0^z \frac{dz'}{H(z')},$$

where $H(z) = \dot{a}/a$ is the Hubble parameter.

Furthermore, all frequencies are shifted down by a factor $1 + z$. For compact binary coalescences, this can be simply accounted for by changing the reference scale of the system, which is encoded in the masses: the signals we measure are then parameterized in terms of *detector-frame* masses M_d , which are related to those in the *source frame*, M_s , by

$$M_d = (1 + z)M_s.$$

For signals emitted by black holes, this is the only measurable combination through gravitational waves: changing z and M_s in a way that leaves M_d unchanged gives rise to exactly identical signals.

6.1.3.1. Gravitational wave cosmology

A significant issue in modern cosmology is the precise determination of the expansion rate of the Universe as a function of time, or equivalently the Hubble rate $H(z)$ as a function of redshift z .

Measurements of distance associated with redshift in the form $(d_L, z)_i$ would be sufficient to uniquely determine it. However, the observation of gravitational waves from a black hole binary only provides estimates of $(d_L, (1 + z)\mathcal{M})$. Much of the field of gravitational wave cosmology is devoted to the disentanglement of these contributions. There are several proposed methods for this [3], some of which have already provided constraints on the local Hubble rate H_0 .

- Source galaxy association, i.e. *bright sirens*: if the sky position of the source is measured well enough that an electromagnetic counterpart is detected and associated with a galaxy, we can use standard spectroscopic methods to determine the redshift z of that galaxy and associate it to the source. This was done for GW170817 [4], the only gravitational wave source with an electromagnetic counterpart to date.
- Spectral methods based on the mass distribution, i.e. *dark sirens*: even if we do not know the mass of any specific source, we can model the mass distribution of the population to extract information on it.
- Methods based on matter effects, i.e. *Love sirens*: if the tidal polarizability parameter Λ of a neutron star (see Section 5.4) is measured, it can be used to infer a mass scale independently from the overall amplitude and phase of the signal.

6.1.3.2. Priors

In an Euclidean Universe, the requirement of the distance prior being volumetrically uniform would correspond to $\pi(d) \propto d^2$.

In an expanding Universe, the equivalent requirement is for the prior to be uniform in comoving volume as well as source frame time; this leads to [151]:

$$\pi(d_L) = \frac{1}{1+z} \frac{dV_c}{dz} \frac{dz}{dd_L}. \quad (6.5)$$




Figure 6.2.: Cosmological priors for luminosity distance. The labels in the legend denote the quantity each prior is uniform over. All priors have been normalized so that they integrate to 1 over the range shown.

In Figure 6.2 we visually compare some options for cosmological distance priors:

- the Euclidean “luminosity volume” prior $\pi(d_L) \propto d_L^2$;
- the “comoving frame volume” prior, uniform in dV_c/dd_L ;
- the “source frame volume” prior, given in Equation 6.5; this is the distribution which we would expect binary mergers to follow in the Universe if the merger rate were constant as a function of time.

6.1.4. Timing

The timing parameter t_0 has a straightforward physical meaning, encoding our uncertainty on the precise moment a given signal has arrived at a reference position. For short signals this is indeed the case, and the timing parameter’s effect can be straightforwardly implemented. For long observations, though, some more care needs to be taken.

Priors on this parameter are generally chosen to be uniform.

6.1.4.1. Detector motion and reference frame choice

In order to quantify the effect of detector motion on the signal, we need to describe the relative motion of source and detector. We will focus on the case of extragalactic signals, as those encompass the totality of confirmed gravitational wave observations so far; galactic sources would require a slightly different but analogous treatment.

It is useful to think of their relative motion as a composition of relative displacements as follows:

$$\text{source} \rightarrow \text{galaxy} \rightarrow \text{Milky Way} \rightarrow \text{SSB} \rightarrow (\text{geocenter}) \rightarrow \text{detector},$$

where by “galaxy” we refer to the source’s galaxy. The center of the Earth (geocenter) is included here for completeness, as it is standard to adopt it as a reference frame center in the analysis of ground-based data, but the barycenter of the Solar System (SSB) is used directly in some cases, bypassing it.

It is possible, though unlikely, that the motion of the source itself is measurably accelerated compared to its environment or its galaxy [175], but even in this case, we do not have ephemeris available to account for it, and any such effects need to be modeled together with other intrinsic source parameters.¹ The motions of the Milky Way galaxy and of the Solar System within it all have very low accelerations: for the purposes of gravitational wave observation lasting at most a few years, the barycenter of the Solar System can be considered to be in uniform motion compared to the source, with an effective velocity vector

$$\vec{v}_{\text{eff}} = \vec{v}_{\text{source} \rightarrow \text{galaxy}} + \vec{v}_{\text{galaxy} \rightarrow \text{MW}} + \vec{v}_{\text{MW} \rightarrow \text{SSB}} + \vec{v}_{\text{SSB} \rightarrow \text{geo}}.$$

Using an effective velocity vector $\vec{v}_{\text{galaxy} \rightarrow \text{MW}}$ is improper in a general-relativistic sense, but for the purposes of data analysis a signal affected by cosmological redshift is indistinguishable from one affected by a Doppler shift.

Let us compare the magnitude of these vectors:

1. $\vec{v}_{\text{source} \rightarrow \text{galaxy}}$, the peculiar velocity of the source, is generally unknown and hard to model. As we shall discuss, gravitational observations alone offer little help in this regard.

¹For a source subject to strong gravitational lensing, the Doppler shift due to velocity alone as opposed to acceleration may be measurable. The point above still applies, though: the velocity and lens parameters are unknown.

2. the magnitude of $\vec{v}_{\text{galaxy} \rightarrow \text{MW}}$ is related to the redshift as $1+z = \gamma(1+|\vec{v}_{\text{galaxy} \rightarrow \text{MW}}|/c)$, while its direction is radially outward. For the gravitational wave observations performed so far, this has ranged from $z \approx 10^{-2}$ [4], corresponding to $v \sim 10^{-2}c$, to $z \gtrsim 1$ [177], corresponding to $v \gtrsim 0.6c$.
3. $\vec{v}_{\text{MW} \rightarrow \text{SSB}}$ is on the order of $2 \times 10^5 \text{m/s}$, less than $10^{-3}c$.
4. $\vec{v}_{\text{SSB} \rightarrow \text{geo}}$ is on the order of $3 \times 10^4 \text{m/s}$, less than $10^{-4}c$.

For typical CBC sources, the cosmological contribution is completely dominant among these effects. So, we can find a reference point² within the Solar System, whose motion can be considered to have constant velocity during the observation period to a sufficiently good approximation. Then, we can assume that this point receives a signal unaffected by Doppler modulations beyond a constant shift in frequency in the form $f \rightarrow f/(1+z)$. Typically, this point is chosen to be the geocenter for ground-based observations of compact binary coalescences, and the SSB for longer observations, such as those of continuous wave signals from spinning neutron stars, or observations with the Laser Interferometer Space Antenna (LISA). In Section 8.3 we will explore the options available in the context of the Lunar Gravitational Wave Antenna.

Once the point has been fixed, it is straightforward to compute the time-dependent delay

$$\Delta t_d(t_{\text{det}}) = \frac{\hat{k} \cdot \vec{r}_d(t_{\text{det}})}{c},$$

where $\hat{k}$ is the unit vector aligned with the signal's propagation direction, while $\vec{r}_d(t)$ is the detector's position vector.

This defines the following mapping between time as measured at the detector, t_{det} , and the time parameter t used to compute the theoretical gravitational wave strain:

$$t_{\text{det}} = t + t_0 + \Delta t_d(t_{\text{det}}),$$

where the constant reference time t_0 is chosen so that the reference frame center receives the event corresponding to $t = 0$ in the waveform at $t_{\text{det}} = t_0$.

For short observations, which includes all the signals measured so far by the LVK collaboration, the time dependence of Δt_d is negligible.

6.2. Stationary phase approximation

The stationary phase approximation [43] (SPA) is a technique to quickly compute an approximate Fourier transform of slowly-modulated oscillatory functions with monotonically

²By “point” here we mean both a point and its motion over time: in the language of General Relativity, a worldline.

increasing frequency. Given a time domain function $g(t) = A(t) \exp(-i\Phi(t))$, its Fourier transform is approximately

$$\tilde{g}(f) \approx A(t^*(f)) \exp(-i\Phi(t^*(f)) + 2\pi f t^*(f)) \sqrt{\frac{-2\pi i}{\ddot{\Phi}(t^*(f))}},$$

where the stationary time $t^*(f)$ is defined by the condition $2\pi f = \dot{\Phi}(t^*(f))$, and dots denote derivatives with respect to time.

It is often useful to compute the stationary time based on the frequency domain phase,

$$\angle \tilde{g}(f) = -\phi(f) = -\Phi(t^*(f)) + 2\pi f t^*(f) - \frac{\pi}{4},$$

where $\angle z$ denotes the angle of a complex number z .

By differentiating it we get a commonly used expression to recover the SPA time in terms of the frequency-domain phasing [120, eq. 21]:

$$\frac{d\phi}{df} = -2\pi t^*. \quad (6.6)$$

The conditions required for this approximation to hold are [43, 180]:

1. $\dot{\Phi}$ is monotonic: $\ddot{\Phi} > 0$;
2. the amplitude $A(t)$ evolves slowly: $|\dot{A}/A| \ll |\dot{\Phi}|$;
3. the frequency $\dot{\Phi}$ evolves slowly: $|\ddot{\Phi}| \ll |\dot{\Phi}|^2$.

6.2.1. SPA as linearization in frequency

The SPA can be viewed as a linearization in phase of the sidebands arising due to a slow modulation. We show this with an example discussed by Chen and Johnson-McDaniel [46].

Let us consider the case of a signal $h(t)$, with known Fourier transform $\tilde{h}(f)$, modulated by a sinusoid as $h_s(t) = \sin(2\pi F t) h(t)$, with $F \ll f$. This is a simplified model for the effect of the detector's orbit on a long-lived, high frequency signal.

The exact Fourier transform of the modulated signal includes contributions from two sidebands to either side of the signal:

$$\tilde{h}_s(f) = \frac{\tilde{h}(f + F) - \tilde{h}(f - F)}{2i}.$$

The SPA result, $\tilde{h}_s^{\text{SPA}}(f) = \sin(2\pi F t^*(f)) \tilde{h}(f)$, can also be recovered by Taylor expanding the phase $\Psi = \angle \tilde{h}$ in the $f \pm F$ terms around f . Since the amplitude is slowly varying, we take $|\tilde{h}(f)| \approx |\tilde{h}(f \pm F)|$. Then:

$$\begin{aligned} \tilde{h}_s(f) &\approx \frac{|\tilde{h}(f)|}{2i} [\exp(i\Psi(f) + i\Psi'(f)F) - \exp(i\Psi(f) - i\Psi'(f)F)] \\ &= \sin(\Psi'(f)F) \tilde{h}_s^{\text{SPA}}(f). \end{aligned}$$

6.3. Calibration

When analyzing data from a detector, we need to estimate strain based on some other physical quantity, such as the number of photons arriving at a photodiode or the displacement measured by a seismometer.

For ground-based interferometers, various data streams (*e.g.* the residual test mass displacement and the analog control displacement) are used to obtain a point estimate for the time-domain strain $h(t)$. Building and continuously validating this process is known as detector **calibration**, and the residual uncertainty in the amplitude and phase of the measured strain is called **calibration uncertainty**.

Calibration uncertainty is commonly modelled as a frequency-dependent deviation around the observed waveform, in the form³

$$h_{\text{obs}}(f) = h_{\text{true}}(f)(1 + \delta A(f)) \exp(i\delta\phi(f)).$$

It is common practice to model the functions $\delta A(f)$ and $\delta\phi(f)$ through a Gaussian Process [38], whose parameters are inferred together with those of the gravitational wave source.

Over the first observing runs for the LVK collaboration, the typical order of magnitude for the 1σ width of the posterior predictive distribution on the calibration envelope was $\delta A \sim \text{few} \times 10^{-2}$ and $\delta\phi \sim \text{few deg}$.

For the GW250114 event, the strain was well-calibrated at both Hanford and Livingston, requiring minimal corrections. This was practically useful, as it enabled a reliable simplified ringdown analysis without accounting for the impact of calibration. We verified this by comparing an analysis accounting for calibration uncertainty with one that assumed the calibration was perfect; the systematic difference between these was below the level due to other systematics such as a different waveform approximant [2].

³The phase term can be substituted by an expression in the form $(2 + i\delta\phi(f))/(2 - i\delta\phi(f))$, which is equivalent to third order in $\delta\phi$ and does not involve exponentials [74].

Based on publicly available data [113], in the left panel of figure Figure 6.3 we show the difference in all parameter posteriors between these two analyses, both of which used the NRSur7dq4 waveform model [187].




Figure 6.3.: Cumulative distributions for all sampled waveform parameters, comparing an analysis marginalized over calibration uncertainties and one which neglected them (*i.e.* setting all the calibration corrections to zero). The bias due to not including calibration is small for GW250114, and significantly higher for GW250207. This figure uses publically available data from the respective data releases [50, 113].

The right panel concerns another event: GW250207, which was detected a few weeks after GW250114; with an SNR of approximately 69, it placed as the second-clearest signal to date[115], its properties (component masses, spins) were quite similar to those of GW250114. While the LIGO Livingston, LIGO Hanford and Virgo interferometers were all taking data, Hanford was not yet in observing mode: hence, its data was not well-calibrated. This is illustrated in Figure 6.3 by the significant difference between the posterior distributions obtained when not marginalizing over calibration uncertainty.

Typically, information about detector calibration comes from *in-situ* measurements; for the first time, however, the astrophysical signal GW250207 was loud enough to significantly inform the calibration of the Hanford detector.

We show this in the right panel of Figure 6.4, where we compare the posterior predictive distribution on the calibration envelope for the Hanford detector in two cases: either an




Figure 6.4.: Calibration envelopes for GW250114 and GW250207 at the LIGO Hanford detector. All posterior envelopes are showing the a 68% symmetric interval. The 68% width of the prior envelopes adopted is also shown with horizontal lines. Adapted from figure 2 of LIGO Scientific, Virgo, and KAGRA Collaborations [115], using publically available data from the releases [50, 113].

analysis of data from all three interferometers, or only from the Hanford detector itself. Using data from LIGO Livingston and Virgo allows for a tighter constraint on the calibration of LIGO Hanford.

In the left panel we provide a comparison with the posterior calibration envelope for the GW250114 event. Besides being significantly narrower, it is near zero at all frequencies, indicating that the detector was properly calibrated at that time. This is compatible with the small degree of bias observed when neglecting calibration uncertainty (Figure 6.3) is minimal for GW250114, while it is significant for GW250207.

6.4. Antenna patterns

Antenna patterns allow us to compute the scalar strain measured by a detector, $h(t)$, starting from the gravitational wave strain tensor $h_{ij}(t)$.

This is often quantified by computing the response of a static detector to a monochromatic wave with frequency f_0 . If this response is linear, it can be expressed as a contraction with a **detection tensor** D_{ij} , which depends on the orientation and shape of the detector's arms:

$$h(t; f_0) = h_{ij}(t; f_0) D_{ij}(f_0).$$

If the detector is not static, this expression needs to be made time-dependent:

$$h(t; f_0) = h_{ij}(t; f_0) D_{ij}(t; f_0).$$

In many applications, this dependence is slow enough to be either ignored completely. When this is not the case, it can often be modelled as an amplitude correction directly in the frequency domain through a Stationary Phase Approximation [95].

If the strain tensor is decomposed as in Equation 6.4, we can express the detection tensor in terms of its effect on the two gravitational wave polarizations:

$$F_{+/\times} = D_{ij} e_{ij}^{+/\times}.$$

6.4.1. Interferometer antenna patterns

Let us start with a derivation of the antenna pattern for an interferometric detector in the static limit, also called long-wavelength approximation, in which the gravitational perturbation h_{ij} is taken to be constant [72]. This is approximately valid when the wavelength of the gravitational wave is much longer than the round-trip travel time along the detector arm, $2L/c$ where L is the arm length. Equivalently, it is approximately valid when the frequency f_0 of the gravitational wave is much smaller than the *free spectral range* $f_{\text{fsr}} = c/2L$.

In this limit, the change in arm length can be taken to be $\delta L = L(h_{ij}/2)l_i l_j$, where n is the unit vector along the arm direction. Measuring the relative displacement in arm length then means computing

$$\frac{\delta L}{L} = \frac{\delta L_x - \delta L_y}{L} = \frac{1}{2} \underbrace{(l_i^x l_j^x - l_i^y l_j^y)}_{D_{ij}} h_{ij}.$$

For high-frequency observations performed with next generation detectors, the long wavelength approximation starts to break down. Figure 6.5 illustrates this: already at a few kHz, ignoring this effect would correspond to a bias of tens of degrees in phase for the Einstein Telescope. Essick, Vitale, and Evans [72] have pointed out that, for sources localized by a network of detectors, neglecting this effect would lead to a bias in the estimate of the arrival time at the detectors, which would in turn bias sky localization.

6.4.2. Antenna patterns for the Lunar Gravitational Wave Antenna

While the operational principle of the Lunar Gravitational Wave Antenna is quite different from that of interferometric detectors, its response can ultimately be cast into a similar form. However, unlike interferometers, the displacement measured by a lunar seismometer is not directly proportional to strain in a wide band. Instead, the lunar response is expected to have significant fluctuations as a function of frequency, as we will discuss in more detail in Section 8.1.3. Therefore, it is convenient to define a response tensor $\mathcal{D}_{ij}$ with the dimensions of a length, such that the displacement $s(t; f_0)$ measured by the detector when a monochromatic gravitational wave impacts the Moon is

$$s(t; f_0) = \mathcal{D}_{ij} h_{ij}.$$

This response tensor can be approximated as

$$D_{ij}(f_0) = \hat{b}_i \hat{n}_j L(f_0),$$




Figure 6.5.: Frequency-dependent antenna response as a function of the angle between source and detector arm. Adapted from figure 1 in Essick, Vitale, and Evans [72], implementing their equation 4. The amplitude of the response is symmetric under sign flips for the observation direction, hence the curves from negative values are not visible in the upper panel. The color scale shows the cosine of the angle between the gravitational wave’s propagation direction and the detector arm.

where $\hat{b}$ is a horizontal unit vector along the lunar surface, along whose direction the seismometer is measuring displacement, while $\hat{n}$ is the normal unit vector to the surface, and the function $L(f_0)$ quantifies the lunar response. We will motivate this expression in Section [8.1.3](#).

Part III.

Future gravitational wave detectors

7. Einstein Telescope

In this chapter, we outline the concept for the Einstein Telescope detector, summarizing some main aspects related to the “Blue Book” we recently published [3], especially related to the challenges of accurate parameter estimation.

Then, we discuss in detail an injection performed for a high-SNR GW250114-like signal (Section 7.2.1), which we presented in Tissino et al. [183]. Finally, we present the results of population-scale analyses comparing the Δ and 2L configurations, that we presented in Santoliquido et al. [159] and Santoliquido et al. [155].

7.1. Detector concept

The Einstein Telescope [146] is a planned next-generation (XG) gravitational wave detector in Europe. It will consist of interferometers, like current detectors, but with several improvements enabling a ten-fold improvement in sensitivity [93]:

- it will be built underground, allowing for a significant reduction in seismic and anthropogenic noise;
- it will have a xylophone design: each detector will consist of two colocated interferometers, one operating cryogenically with low laser power to be sensitive at low frequencies, and one operating at room temperature at higher power to be sensitive at high frequencies;
- it will have longer arms, of 10km or more in length depending on the design.

There are currently two main proposed designs. One option is a set of three interferometers, each with an aperture angle of 60 degrees, forming an equilateral triangle. This is called ET- Δ . The alternative consists of two L-shaped detectors, each with an aperture angle of 90 degrees, in two different locations within Europe, with a distance of little more than 1000km. These two detectors may be aligned, like the two LIGO detectors, or misaligned — the choice between these configurations involves, among other considerations, a trade-off between sensitivity to compact binary sources and stochastic backgrounds [29].

There are different sites which may, as of this writing, host the Einstein Telescope within Europe [160]:

- the Sos Enattos mine in the northeast of Sardinia, Italy, near the villages of Bitti, Onanì and Lula — we depict this in Figure 7.1;
- the Euregio Meuse-Rhine (EMR) at the border of Luxembourg, Belgium and Germany, near the towns of Maastricht and Liège;
- the Lusatia region in Saxony, Germany.




Figure 7.1.: Sardinian candidate site for the Einstein Telescope, the configuration used for the simulations discussed in Section 7.2.1.

The scientific potential for the Einstein Telescope is extremely wide, including CBC population studies, precision cosmology, multimessenger observations, and tests of general relativity. A full description is beyond the scope of this thesis, and we refer the reader to some among the many works which describe it in detail [29, 3, 153].

In a similar time frame, a detector dubbed Cosmic Explorer is planned in the United States [148]. While it shares many similarities with the Einstein Telescope, there are some major design differences between the two. Cosmic Explorer’s planned arm length is 40km, significantly longer than any of the designs for the Einstein Telescope, and it will not use cryogenics. Hence, its sensitivity is expected to be better at high frequencies, but worse at low frequencies.

In the remainder of this chapter we shall discuss the challenges which come when trying to perform parameter estimation with a very sensitive detector, and how some recent

advancements in these techniques are allowing for systematic forecasting studies, which can help in informing the urgent choices that need to be made regarding the construction of this instrument. While some of the results presented here are generally applicable to any sensitive ground-based interferometer, we also shall explore some challenges specific to the candidate geometries for the Einstein Telescope.

7.2. Data analysis challenges

This section will give a brief summary of the primary challenges that remain unsolved in the data analysis of compact binary signals with the Einstein Telescope. They are described in significantly more detail in chapters 8 and 10 of Abac et al. [3], which also outlines some problems we will not consider here, such as the detection of CBCs and the analysis of different types of signals like continuous waves or stochastic backgrounds.

The main differences between the analyses performed with the Einstein Telescope, compared to current detectors, are as follows.

- The noise is overall lower, by approximately an order of magnitude, which leads to
 - more detected signals, approximately 10^5 black hole and neutron star binaries per year in total,
 - higher signal-to-noise ratios for the closest events, reaching the level of $\rho \approx 10^3$.
- The sensitive band extends to both higher and lower frequencies than current detectors.
 - At lower frequencies and for low-mass binaries, this leads to them being in-band for several hours, which makes the effect of the Earth’s rotation on the angles between detector and source non-negligible.
 - At higher frequencies, the improved sensitivity combined with the longer arm length leads to the necessity of considering a frequency-dependent antenna pattern as discussed in Section 6.4.

Including time and frequency dependence in the antenna patterns is not an issue *per se*, as the way to account for these effects computationally is well-known. These effects significantly increase the complexity of the space of possible waveforms which may reach the detector, and can in fact be beneficial by making signals with different parameters more distinguishable. However, some of the current approaches for faster parameter estimation significantly suffer from their inclusion [15]: the computational footprint for a reduced order model (Section 3.4.2), which guarantees accuracy for any point in a given prior space, dramatically increases when the antenna patterns are measurably time- and frequency-dependent.

The number of signals being much larger poses some issues of its own, besides the obvious one of requiring a proportionally larger amount of computational resources. With the current generation of detectors, it can generally be assumed that the time before and after any given signal only contains noise, which allows for *off-source* noise characterization. This will not be true anymore, especially in the presence of many long-lived, low-mass sources.

Furthermore, some of these signals may overlap [see *e.g.* 100]; this breaks another common assumption in gravitational wave analysis, that of only one signal being present in the data segment being analyzed.

Finally the high SNRs of the best-measured signals, while being crucial for many parts of the Einstein Telescope’s science case (multimessenger science, tests of GR and so forth) significantly challenges several aspects of the current infrastructure, with waveform modelling likely being the most affected. We can draw an approximate relation, based on a Fisher matrix argument, between the maximum mismatch (Equation 5.4) which is expected not to bias parameter estimation at a given SNR ρ [14]:

$$\bar{\mathcal{F}} \lesssim \frac{\chi_D^2(1-p)}{2\rho^2},$$

where $\chi_D^2(p)$ denotes the inverse CDF of a chi-square variable with D degrees of freedom, where $D \approx 15$ is the number of waveform model parameters.

For example, at a 90% credible level with 15 parameters, the constraint is $\bar{\mathcal{F}} \lesssim 11.15/\rho^2$; at an SNR of 1000, this translates to $\bar{\mathcal{F}} \lesssim 1.1 \times 10^{-5}$.

This level of accuracy is currently beyond the reach of any waveform model; the situation is better for low spin and eccentricity, but outside of this regime biases are already observable with *current* detectors [1].

Overall, significant work remains to be done in the next decade in order to prepare for the fast and unbiased analysis of Einstein Telescope data.

7.2.1. GW250114 as seen by a triangular Einstein Telescope

We inject a waveform corresponding to the maximum likelihood point from the LVK analysis of the gravitational wave event GW250114 into the data of a triangular Einstein Telescope. We analyze the data using the software *bilby* [11] combined with the sampler *nessai* [195]. Our analysis does not include time- nor frequency-dependent antenna patterns. We summarize the priors we adopt, as well as the injected values for various parameter, in Table 7.1.

Table 7.1.: Priors and injected values for our GW250114-like injection with the Einstein Telescope. The timing parameter Δt represents a deviation from the injected value, a GPS time for the merger of 1420878141.24 seconds, as measured at the geocenter.

Parameter	Injected value	Unit	Prior function	Prior range
$\mathcal{M}$	31.3	$M_{\odot}$	Uniform	[30, 32]
q	0.98		Uniform	[0.95, 0.99]
χ_1	-0.05		Aligned spin (Equation 5.2)	[-0.9, 0.9]
χ_2	0.01		Aligned spin (Equation 5.2)	[-0.9, 0.9]
d_L	414	Mpc	Uniform in the source frame (Equation 6.5)	[400, 450]
θ_{JN}	0.7	rad	Sine	[0, π]
$\Delta t_{\text{merger}}^{\text{geocenter}}$	0	s	Uniform	[-0.10, 0.10]
RA	2.3	rad	Uniform	[0, 2π]
DEC	0.2	rad	Cosine	$[-\pi/2, \pi/2]$
ψ	1.3	rad	Uniform	[0, π]
ϕ	1.6	rad	Uniform	[0, 2π]
m_1	36.3	$M_{\odot}$	Constraint	[35, 37]
m_2	35.5	$M_{\odot}$	Constraint	[34, 36]

This signal would have an SNR of approximately 680 when combining the three detectors. Accordingly, the constraints on the source’s intrinsic parameters are quite tight; we shall show some of them in Section 8.2.3.3, with a comparison to the Lunar Gravitational Wave Antenna.

Here, instead, we shall focus on a peculiar feature of the posterior distribution. In Section 7.3.1 we describe how short signals are typically expected to exhibit an 8-fold multimodality in the sky. We only observe two modes in the sky for this injection, symmetrically across the local horizon, as shown in Figure 7.2. Specifically, about 5% of the posterior mass is in this mode, which also features a significantly different geocentric time, as well as inclination and polarization angles peaking at π minus their injected value.

The antenna pattern symmetry leading to this mode can be understood analytically; see also equation 63 in Marsat, Baker, and Canton [121]. Why the posterior distribution is not perfectly symmetric, and why the other 6 modes are not present, are not clear from the discussion in that work, as for this injection we employed the long wavelength approximation (LWA).




Figure 7.2.: Posterior distribution on the sky position for our GW250114 injection. We show individual posterior samples with a scatter plot, and color them by the corresponding estimated geocentric time, quantified as a deviation from its injected value. We mark the injected position with a white cross.

What those authors do not discuss (as it is not relevant in the LISA case) are the features of posterior distributions obtained in the LWA, but accounting for the separation between beamsplitters. In our simulation, the Einstein Telescope consists of three detectors separated by 10km each which, due to the long-wavelength approximation, are effectively point-like. The temporal resolution for the arrival time of the signal at each them is so precise, on the scale of tens of microseconds, that it allows (our simulation of) the ET to *triangulate* the signal, despite its extremely short baseline of $33\mu\text{s} \times c$. In order to show this, in Figure 7.3 we show the posterior distribution on the arrival time of the merger at the three beamsplitters. The three arrival times are well-measured on the scale of tens of microseconds, thus the time delays between detectors can inform the sky localization.




Figure 7.3.: Posterior distribution on the arrival times at the three ET beamsplitters in our GW250114 injection. From Tissino et al. [183].

7.3. Neural posterior estimation for high-mass black holes

In order to investigate the possible configurations for the Einstein Telescope, we [159, 158] performed several injection campaigns with neural posterior estimation (Section 2.3.2), which provide full parameter estimation results for each of the sources. With the current technology, this is possible for a population of high-mass black holes at high redshift, which last for a short time in the detector band and have generally low SNRs. Though we inject sources in the chirp mass and distance ranges $\mathcal{M} \in [40, 1100]M_\odot$ and $d_L \in [5, 500]\text{Gpc}$, we are forced to discard most sources with chirp masses $\mathcal{M} \lesssim 150M_\odot$ and distances $d_L \lesssim 6\text{Gpc}$, due to their low sample efficiency, shown in Figure 7.4. Specifically, we discard any source for which the sample efficiency, defined in Section 2.3.3, is lower than 10^{-2} .

Sources similar to these are expected to constitute the majority of black holes observed




Figure 7.4.: Sample efficiency for the sources considered by Santoliquido et al. [158] as a function of SNR, for all detector configurations. We observe little variation across configurations. For SNRs higher than 100, we generally observe low sample efficiencies: as long as that is the case, the usage of NPE does not present a computational advantage for these sources compared to stochastic methods.

by the Einstein Telescope. Black holes are expected to form at high redshift from several processes, among them the collapse of the nearly metal-free first generation of stars (Population III). Santoliquido et al. [156] investigated their formation history, highlighting the uncertainties in our understanding of these early stars. Constraining the properties of these black holes thus has interesting astrophysical implications.

While an investigation including only a specific kind of source can only provide partial information, it can still highlight some important characteristics of each detector network's capabilities.




Figure 7.5.: Summary statistics for source position constraints by various configurations. The information gain (Section 2.2.6) encodes the global volumetric compression between prior and posterior. We see variations of one to two orders of magnitude across configurations for the sky localization area, fractional distance uncertainty and comoving volume constraints. This is driven by network geometry, as we can see from the fact that the distribution of SNRs is similar for all configurations. From Santoliquido et al. [158].

In Figure 7.5 we provide some summary statistics for several configurations of Einstein Telescope, also combining it with other detectors.

We will focus specifically on three configurations, which involve the Einstein telescope alone:

- ET - Δ , with 10 km arms, which we will discuss in more detail in Section 7.3.1;
- ET - 2L, with 15 km arms, in two configurations: aligned and misaligned by 45 degrees, discussed in Section 7.3.2. The two detectors are assumed to be in the EMR and Sardinia respectively. The aligned and misaligned configurations are labelled “2L A” and “2L MisA” respectively.

The comparison also includes other options with additional detectors, as follows:

- 2L MisA + LHI involves two misaligned L-shaped Einstein Telescope detectors, as in the “ET-2L MisA” configuration, paired with a network of the LIGO detectors at Hanford and Livingston plus LIGO India, with all three latter operating at their mid-2030s “A#” sensitivity;
- 1L + CE: an L-shaped Einstein telescope in Sardinia combined with a 40 km Cosmic Explorer located near Hanford;
- Δ + CE: a triangular Einstein Telescope in Sardinia combined with a 40 km Cosmic Explorer located near Hanford;
- 2L MisA + CE: two Einstein Telescopes in Sardinia and the EMR region, each with 15 km arms, combined with a 40 km Cosmic Explorer located near Hanford.

Including a distant detector improves the baseline of the measurement, hence these configurations are typically better at localizing than Einstein Telescope alone. In any case, the focus of the following discussion will be a comparison of Einstein Telescope designs.

For the analysis of the events in our sample, the misaligned 2L configuration is generally more informative about the signals, though the Δ configuration typically constrains their distance more precisely. The misaligned 2L configuration is better at constraining sky position, which translates to smaller values of the comoving localization volume.

We started from the extrinsic parameters — distance and sky position — as they are the most dependent on detector configuration. As we see in Figure 7.6, the errors in most intrinsic parameters obtained with different configurations are quite stable, as they depend on the source SNR alone. This can also be understood from the correlation structure of the posterior distributions: see Section 2.4.1.3.

The only parameter for which we see significant differences in Figure 7.6 is the inclination angle, which depends on the amplitudes observed by different detectors, and is therefore generally correlated with luminosity distance.




Figure 7.6.: Summary statistics for intrinsic parameter constraints by various configurations. From Santoliquido et al. [158].

7.3.1. Δ shape

The Δ design for the Einstein Telescope consists of three co-located interferometers, the arms of each are two sides of an equilateral triangle. These three interferometers are labelled ET-1, ET-2 and ET-3. This detector design has the unique property of being able to “self-localize”, that is, determine the position of a gravitational wave source with high precision without making use of travel-time measurements between detectors.

For the many CBC sources, though, this precise localization is sullied by an eight-fold multimodality [159], which was also previously highlighted in the context of LISA [121], which shares the triangular geometry. An example of this multimodality is shown in Figure 7.7.




Figure 7.7.: Sky position posterior, colored by arrival time at geocenter, for a massive black hole injection as observed by a triangular Einstein Telescope. The posterior is shown in an altitude-azimuth coordinate frame centered at the Einstein Telescope’s location. From Santoliquido et al. [159].

These degeneracies are driven by symmetries in the antenna patterns, and they are expected to be present as long as no effect is able to make some modes preferred over others. The following effects could do so:

- Modulation driven by the Earth’s rotation for a long signal. The precise threshold for this will depend on many specifics, but we can expect it to be on the scale of an hour in duration.

- Observation with another detector, not co-located with the triangular Einstein Telescope.
- Very high SNRs, which make the non-point-like nature of the detector appreciable, as in Section 7.2.1, though this aspect remains under investigation.

We now turn to the geometric reasoning behind the multimodalities. In terms of the local azimuth λ , altitude β , and the polarization angle ψ of the incoming gravitational wave, the antenna pattern functions read [147]:¹

$$F_+^{\text{ET-1}}(\beta, \lambda, \psi) = -\frac{\sqrt{3}}{4}[(1 + \cos^2 \beta) \sin 2\lambda \cos 2\psi + 2 \cos \beta \cos 2\lambda \sin 2\psi]$$

$$F_\times^{\text{ET-1}}(\beta, \lambda, \psi) = +\frac{\sqrt{3}}{4}[(1 + \cos^2 \beta) \sin 2\lambda \sin 2\psi + 2 \cos \beta \cos 2\lambda \cos 2\psi].$$

The antenna patterns for ET-2 and ET-3 are obtained with the transformation $\lambda \rightarrow \lambda \pm 2\pi/3$.

These are all symmetric under a reflection across the local horizon ($\beta \rightarrow -\beta$), rotation around the locally vertical direction by half a turn ($\lambda \rightarrow \lambda + \pi$) and simultaneous rotation of the azimuth and polarization angle by a quarter of a turn ($\lambda \rightarrow \lambda + \pi/2, \psi \rightarrow \psi + \pi/2$).

These symmetries can be combined in any way, leading to an eight-fold degeneracy in the sky position of any source, with four equally-spaced modes above the local horizon, and four below it.

This also leads to a bimodality in time at the geocenter, if that is used as a parameter. Arrival time at the detector is directly measured, and to convert it to time at the geocenter we must add a term $\hat{n} \cdot \vec{r}_{\text{ET}}/c$, which also equals $(|\vec{r}_{\text{ET}}|/c) \cos \beta$. The bimodality in β , then, is propagated to geocentric time. This can be seen in Figure 7.7, which shows the posterior distribution on the sky position in local altitude/azimuth angular coordinates colored by geocentric time.

In Figure 7.8 we summarize the number of modes in the sky location posterior distribution for the various configurations considered. A noticeable fraction of sources exhibits eight-fold multimodality with the Δ configuration. Not all of them do, since in some cases the posterior may be broad enough not to qualify as “multimodal” with the clustering algorithm we are using even if it is symmetric; alternatively, for sources close to the local horizon the modes with positive and negative β may not be distinguishable, leading to a 4-fold multimodality.

¹In Santoliquido et al. [159] the labels “altitude” and “azimuth” are inverted when defining these antenna patterns.




Figure 7.8.: Number of modes in the sky for a selection of Einstein Telescope configurations and other detectors. From Santoliquido et al. [158].

The effect we discussed in Section 7.2.1, on the other hand, with the travel time between beamsplitters enabling the exclusion of certain modes, is not observed at the lower SNRs considered here.

7.3.2. 2L shape

While the 2L configuration does not exhibit the eight-fold multimodalities we see with the Δ , it is not free from them, as is apparent from Figure 7.8.

The baseline between the two sites is $|\Delta\vec{r}| = 1166\text{km} \approx 3.9\text{ms} \times c$. If $\Delta\vec{r} \cdot \hat{m}$, where $\hat{m}$ is the source direction, is significantly longer than the uncertainty on the arrival time of the signal at our detectors σ_t , the time delay can inform the sky localization of the source. If this is not the case, the only information available to this end comes from comparing the amplitude of the observed waveform at the various detectors. For the high-mass black holes in our sample, the timing uncertainties are generally large, and therefore we typically find ourselves in this scenario.

The amplitude of the signal, however, is also affected by its distance and inclination angle. This can lead to correlated multimodalities in the joint posterior distribution of sky position and distance, as illustrated in Figure 7.9.

The example in Figure 7.9 also helps in building an intuition for the interpretation of multimodalities. It is tempting, when presented with an asymmetric multimodal posterior distribution, to assume that the true value will lie in the higher mode. Of course, as long as our analysis is unbiased, the probability mass contained in each mode is precisely the probability that the true value will be contained in that mode. In this example, the true value is contained in the less-probable mode.

Multimodalities in distance are observed for a significant fraction of the sources in our sample, for all configurations except for the Δ , as shown in Figure 7.10.

7.3.3. Statistical validation

In order to verify that our results are robust, we perform some statistical consistency tests.

The first among them is a p - p test (Section 2.1.2), shown in Figure 7.11. Furthermore, we perform full parameter estimation with nested sampling (Section 2.2) for a selection of sources, and directly compare the posterior distribution to that obtained with neural posterior estimation.




Figure 7.9.: Sky and distance posteriors for an injection in the 2L, misaligned case, compared to a triangular Einstein Telescope. We show results obtained with stochastic sampling (“Bilby”): they perfectly overlap with those obtained with NPE. The 2L configuration, in this case, yields a multimodal posterior distribution for both distance and sky position. From Santoliquido et al. [158].




Figure 7.10.: Multimodality in the distance posteriors for all configurations. From Santoliquido et al. [158].

We measure their differences through the Jensen-Shannon Divergence, which is defined as a symmetrized version of the Kullback-Leibler divergence (Section 2.2.6):

$$\text{JSD}(p, q) = \frac{1}{2} (\text{KL}(p||q) + \text{KL}(q||p)) .$$

This distance ranges in value from 0 to $\log_b 2$, where b is the base in which the logarithms in the KL divergence are computed.

The JSD between two sets of sample is commonly used to assess whether they originate from the same distribution. In order to do so, we need to determine what values of estimated JSD are compatible with statistical fluctuations, *i.e.* their null distribution. We determine this in a data-driven way: since we can easily sample from the NPE distribution, we draw 100 realizations of the full posterior and compare them with the JSD, in all $100 \times (100 - 1)/2 = 4950$ combinations. This gives us a distribution of Jensen-Shannon divergences under the null hypothesis, which we can compare to the observed divergence with the nested sampling run.

The test confirms that the two distributions are statistically consistent [158]; as far as we are able to check, the inference we are performing is unbiased.




Figure 7.11.: Results of a pp-plot test for all configurations considered by Santoliquido et al. [158]. For all parameters and all configurations, we confirm the statistical self-consistency of our approach.

8. Lunar Gravitational Wave Antenna

In this chapter we introduce the Lunar Gravitational Wave Antenna, following the whitepaper we recently published as a collaboration [6]. Then, we outline some novel contributions, namely the first discussion of the issues related to full Bayesian parameter estimation for compact binaries with this detector, which we presented in Tissino et al. [183] and Iacovelli et al. [96].

8.1. Detector concept

The idea of using a resonant body as a gravitational wave detector was first proposed by Joseph Weber [191], following investigations on the effects of gravitational waves on matter by Pirani [142].

Weber led the efforts to deploy a gravimeter on the Moon with the Apollo 17 mission, with the stated aim to detect gravitational waves. The Lunar Surface Gravimeter [84] aimed for a sensitivity of around 10^{-10} m in vertical displacement when measuring vertical motion, around the 1.5Hz band.

We now know that, even operating nominally, this would have been far too noisy to actually measure gravitational waves. Furthermore, the sensor beam could not be stabilized upon deployment, due to some errors by the manufacturers when converting from terrestrial to lunar gravity.

Nevertheless, the Apollo program provided valuable insights into the structure of the Moon. In the following decades, both technology and our understanding of gravitational waves have advanced.

Harms et al. [89] proposed the concept for a Lunar Gravitational Wave Antenna, which uses seismometers to measure vibrations of the lunar surface induced by gravitational waves. The plan is to deploy this antenna in a Permanently Shadowed Region at the lunar pole (see Section 8.1.1), in order to achieve low-noise cryogenic operation with a displacement noise ASD (Equation 3.2) below one fm/ $\sqrt{\text{Hz}}$. Combined with a sufficiently strong lunar response to gravitational waves, this would be enough to directly measure gravitational waves in the deci-Hertz band.




Figure 8.1.: The Lunar Surface Gravimeter, by NASA Apollo 17 photograph AS17-134-20501, Public Domain.

In the following years, a collaboration has been formed to study the science of the Lunar Gravitational Wave Antenna; this has resulted in the publication of several works, among them the *whitepaper* [6], which details several aspects of the possible science achievable in this band.

This chapter will focus on the full LGWA mission and its projected sensitivity, but a preliminary mission, dubbed Soundcheck, is also planned. Whereas the LGWA will require advancements in sensor technology compared to the current state of the art, Soundcheck will use off-the-shelf components. Even so, if deployed at the lunar pole, it will provide extremely valuable data, though the chance it will detect gravitational waves is slim. The Soundcheck mission has been included in the Reserve Pool of Activities by the European Space Agency.

8.1.1. Site selection

Permanently Shadowed Regions (PSRs) exist in the lunar polar regions: due to the low tilt of the Moon's orbit, the Sun remains close to the horizon throughout the year, therefore some low-elevation regions never receive direct sunlight. While some PSRs receive reflected sunlight, such as from a crater's rim, temperatures within them are typically extremely cold, in the range of tens of degrees Kelvin.

A preliminary site selection study for the LGWA was performed in 2025 by a team at ETH Zürich [85]. The choice of a site crucially depends on power generation: the main options are a radioisotope thermoelectric generator (RTG) and solar power. With the former option, a lander could directly target the deployment region; with the latter, a solar station would need to transmit power to the stations, for example with a long cable.

Gobe et al. [85] identified a potential landing site in the lunar north polar region, at latitude 88.389 and longitude 280.76 degrees, which could be used to then deploy seismometers in a nearby PSR, around the latitude 88.50 and longitude 290.63 degrees. This PSR lies roughly between the Hermite A and Hinshelwood craters. We show these locations in figure Figure 8.2; useful free and interactive mapping tools to explore them include [LROC Northern Polar Mosaic](#) and [Quickmap](#).

The distance between the landing site and the PSR center is only around 8.6km, but their temperature profiles are significantly different. The landing site is a flat hilltop, illuminated for most of the (northern) summer months, April-September. In the winter, it receives at least a week of illumination every month, but it is predominantly in the shade, with continuous darkness for up to 22.4 days. Its temperatures range between 66K and 250K across the year. The PSR, on the other hand, has a mean temperature of 46K, with maximum recorded values of less than 60K. A comparatively flat path connects the landing




Figure 8.2.: LGWA proposed landing site on the North Pole (right blue dot) and PSR (left blue dot), in 3D perspective. The color on the lunar surface indicates maximum summer temperatures, ranging from 50K (deep blue), through 100K (white) to red (150K or more). The Earth's and Sun's apparent orbits are shown for reference. The large crater whose rim is visible to the right is Hermite A. This figure was generated with QuickMap [\[117\]](#).

and deployment sites, with slopes of less than 15 degrees, though the elevation difference between them is significant.

An alternative landing site near the South Pole was also explored, within the de Gerlache crater at latitude -88.381 and longitude 266.8. Both this deployment site and the North Pole sites are rather flat themselves, with maximum slopes of few degrees. However, the South Pole is generally more rugged: the path from an illuminated landing site to the PSR location, which would likely have to be taken by a rover, nears slopes of 20 degrees.

8.1.2. Measurement principle and displacement noise

The operational principle for the LGWA is the measurement of horizontal seismic motion of the lunar surface, which will be performed along two orthogonal axes. This measurement requires a reference whose movement is inertial along the two measurement directions, whereas the surface will move non-inertially, due to the Moon being an elastic body.

This reference will be a mass suspended through a Watt's linkage [22], a compact way to achieve a low resonance frequency.

i Note 11: Watt's linkage

A Watt's linkage is, at its core, the combination of a pendulum and an inverted pendulum.




Figure 8.3.: Watt's linkage diagram, from Bertolini et al. [22].

The resonance angular velocity of the system depicted in Figure 8.3 is

$$\omega_0^{\text{Watt's linkage}} = \sqrt{\left(\frac{M_1}{L_1} - \frac{M_2}{L_2}\right) \frac{g}{(M_1 + M_2)}} + \gamma,$$

where g is the gravitational acceleration, M_i and L_i are the two masses and lengths shown in the figure, while γ represents the effect of cumulative flex-joint stiffness. This allows for very low resonance frequencies with a comparatively small apparatus. For comparison, let us consider a simple pendulum: it achieves the same objective, but its resonance is at

$$\omega_0^{\text{pendulum}} = \sqrt{\frac{g}{L}}.$$

To achieve the resonance angular velocity planned for the LGWA, $\omega_0 \approx 1.57\text{rad/s}$, it would require a length of $L \approx 0.66\text{m}$, making it quite unwieldy.

The formulas above are only valid in the small-angle regime; the plan with the LGWA is to have an actuator keep the linkage in this position. Then, the required output from this actuator will contain the displacement signal.

In Figure 8.4 we show a conceptual overview of a single LGWA station; its subsystems and the tilt of the lunar surface are not to scale. A platform leveling system is crucial in order to avoid vertical-to-horizontal coupling. The expected operating temperatures for the proof masses are much lower than the ambient temperature; this can be achieved through sorption cooling, which involves no mechanical moving parts besides a few passive valves. This is helpful both for the mission's lifetime and for vibration reduction.

The main source of readout noise at low frequencies ($\lesssim 0.7\text{Hz}$) is expected to be suspension thermal noise, which scales as

$$S_{\text{thermal}}(\omega) = \frac{4k_B T \omega_0^2 / Q}{m\omega [(\omega_0^2 - \omega^2)^2 + \omega_0^4 / Q^2]},$$

where ω is the angular frequency, ω_0 is its resonance value, k_B is the Boltzmann constant, T is the temperature, while Q is the mechanical quality factor [6, eq. 2.1].

At higher frequency, readout noise starts to dominate; its specific spectral shape depends on the specific readout scheme adopted to measure relative displacement between ground and test mass. Figure 8.4 schematically depicts a SQUID readout, which makes use of superconducting magnets. In Figure 8.5 we compare this choice with another candidate: laser-interferometric readout.





Figure 8.4.: LGWA station diagram from Joris van Heijningen et al. [102].



(a) Pessimistic model: a Niobium suspension and laser readout.

(b) Optimistic mode: a Silicon suspension and a SQUID readout.

Figure 8.5.: Displacement noise estimates for the LGWA. For both concepts, thermal noise dominates at low frequency, while displacement noise dominates at high frequency.

8.1.2.1. Other noise sources

Displacement sensitivity curves similar to those in Figure 8.5 were adopted in all forecasting studies; no other noise sources were considered.

While a PSR is undoubtedly a quiet place, it is warranted to ask whether the seismic environment of the Moon can interfere with gravitational wave measurements, or prevent them entirely. Unfortunately, the current data on this is quite limited. Coughlin and Harms [54] analyzed one year of data from the Apollo Lunar Surface Experiments Package (ALSEP) mission, which consisted of four seismometers deployed on the near side of the Moon by the Apollo 12, 13, 14 and 16 missions.




Figure 8.6.: Upper bound on the lunar seismic spectrum, from [54]. The dashed lines indicate the minimum and maximum for a model of seismic noise on the Earth.

In Figure 8.6 one can see the average of this measurement. The flat spectrum at around $10^{-10} \text{m}/\sqrt{\text{Hz}}$ is an upper limit, as the instrument’s precision limited the measurement from reaching lower values. Despite this, it is considerably lower than the minimum achieved anywhere on Earth (dash-dotted curve).

From theoretical considerations, we expect the actual seismic background to be significantly lower; furthermore, a spectrum computed this way is not necessarily representative of the “typical” spectrum one would observe at any given time on the Moon: it is a root-mean-square of values observed over a year, which means it includes the contribution of loud events which are well-localized in time. For the purposes of the LGWA, these may simply

lead to a fraction of the data being unusable, as opposed to making the entirety of the data too noisy.

The LGWA will also operate in a PSR without dawns or dusks, removing their thermal stresses as a noise source. Finally, the intra-station distance among the four seismometers is planned to be close enough that the seismic signal from gravitational waves is strongly correlated among them, but far enough that events such as meteorite impacts arrive at each station with measurably different amplitude and phase, and can therefore be subtracted.

Nevertheless, more measurements are needed to quantify the impact of the lunar seismic background. Moonquakes remain a significant issue, and it is likely that some fraction of the data will need to be unusable due to their effect. The low damping of the lunar regolith makes their duration much longer than what typically happens on Earth, as illustrated in Figure 8.7.




Figure 8.7.: Comparison between a moonquake detected within the Apollo 12 mission and an earthquake, from Stein and Wyssession [170]. In both cases, the vertical axis shows seismic displacement due to the quake, with arbitrary normalization.

Other sources of noise yet may affect the measurement, such as human activity from other

nearby projects, magnetic noise from lunar dust and so forth. Mitigating these will be a significant challenge as the mission is developed.

8.1.3. The lunar seismic response

The fundamental lunar normal mode, which describes quadrupolar vibration of the whole Moon, has a frequency just above 1mHz. While plenty of gravitational wave sources are expected at these frequencies, which are a target of the LISA mission, the LGWA detector concept is not expected to be very sensitive there, mostly due to thermal noise (see Figure 8.5). Computing the lunar response at frequencies two to three orders of magnitude higher than the fundamental is challenging with the normal mode formalism, as one needs to include a very large number of modes, and local topography is expected to play a significant role.

The standard treatment for the problem of an elastically coupled medium to gravitational waves was introduced by Dyson [69] and applied to the computation of the excitation of the normal modes of a planetary body by gravitational waves by Ben-Menahem [19]. Recent years have seen significant theoretical and computational advancements in this regard. Belgacem, Maggiore, and Moreau [18] formulated the coupling problem with a fully relativistic Effective Field Theory approach, which generalizes well to high frequencies beyond the reach of the normal-mode approach. Bi and Harms [23] investigated the impact of a layered geological structure using a half-space model, showing that it can significantly amplify the lunar response compared to a uniform Moon. A significant amount of work has also been recently performed regarding numerical simulations of the lunar response [197, 198, 199].

For the remainder of this discussion, we will make two strong assumptions:

1. the lunar response is *linear*, meaning that a gravitational wave at a frequency f_0 only induces a seismic response with the same frequency, without exciting any other vibration;
2. the lunar response has zero phase: the vibration excited is in phase with its source.

Neither of these are going to be exactly realized in practice, and relaxing them is an important part of future efforts. We expect that the zero-phase assumption will be significantly violated in practice, but this is immaterial for the purposes of the studies in this work. As we will discuss in Section 8.2.3.1, we perform injections assuming the response is known, which means that any assumption we could make about the response’s phase would be exactly cancelled within our likelihood.

Under these assumptions, the scalar lunar response L is a function of frequency with units of length, and it is defined as the ratio between the amplitude of the surface displacement

s and the amplitude of the gravitational wave strain h which drives it:

$$L(f_0) = \frac{s(t; f_0)}{h(t; f_0)},$$

where the amplitude of the gravitational wave is obtained by contracting the gravitational wave strain tensor h_{ij} with a geometric detector response $\hat{D}_{ij}$, and where the strain is assumed to be a monochromatic wave with a frequency f_0 .

Here we limit ourselves to a simple justification, within Dyson's description of the problem in the transverse-traceless gauge as outlined by Bi and Harms [23], of the geometry of the detection tensor $\hat{D}_{ij}$. The LGWA sensors will be set up to measure horizontal vibration, as most of the seismic waves near the surface are expected to be vertically propagating shear waves. Within the Dyson description, gravitational waves couple to the elastic medium through the gradient of the shear modulus $\nabla_i \mu(\vec{r})$, giving rise to an effective force term in the form

$$f_i = h_{ij} \nabla_j \mu(\vec{r}).$$

Then, if we want to compute horizontal acceleration along an axis b we will need to consider forces along that direction. Then, the relevant component of the gravitational wave strain tensor will be proportional to

$$h_{ij} b_i \nabla_j \mu.$$

The shear modulus will vary across geologic layers and across the surface: these are horizontally-placed strata, hence the gradient is expected to be vertical. This is the origin of the geometric expression we use in Section 8.1.4; corrections to it will be necessary based on local topography, but for the purposes of current studies the uncertainty in the scalar response is the dominant one.

8.1.4. Projected gravitational wave strain

As we anticipated in Section 6.4.2, the displacement corresponding to an impacting monochromatic gravitational wave can be modeled through a detection tensor

$$\mathcal{D}_{ij}(t; f_0) = n_i(t) b_j(t) L(f_0) = \hat{D}_{ij}(t; f_0) L(f_0),$$

where $n_i(t)$ is the normal to the lunar surface at the detector location at time t , while $b_j(t)$ is a vector parallel to the lunar surface and aligned to the measurement channel. We compute all vectors in an ICRS-aligned frame (where the spherical angles are right ascension and declination), so that the source's position is constant.

In Section 6.1.4 we outlined the effect of detector motion on the waveform in the time domain. Applying a Stationary Phase Approximation, we can combine the effect of a

moving detector and a rotating antenna pattern in the frequency domain: the measured frequency-domain strain will satisfy

$$s(f) = h_{ij}(f) \mathcal{D}_{ij}(t(f); f) \exp \left(2\pi i f \left(t_0 + \frac{\vec{r}(t(f)) \cdot \hat{m}}{c} \right) \right), \quad (8.1)$$

where:

- $t(f)$ is the time-to-frequency map given by the Stationary Phase Approximation (Equation 6.6);
- $h_{ij}(f)$ is the frequency-domain gravitational wave strain tensor;
- $\hat{m}$ is the *propagation* unit vector (which has opposite sign to the *source position* unit vector, and points “inward”);¹
- $\vec{r}(t)$ is the position of the detector as a function of time, measured in an ICRS-aligned coordinate frame, centered at some point in space;
- t_0 is the reference time, *i.e.* the time at which the event which occurs at $t = 0$ in the time-domain waveform $h(t)$ reaches the center of the reference frame, $\vec{r} = 0$.

Let us define an orthonormal reference frame at the lunar surface: (a, b, n) , where n is aligned to the local normal direction, while a and b define the horizontal measurement directions of the two LGWA channel. In order to compute the geometric portion of the antenna pattern tensor, $\hat{D}_{ij}$, we need to know these vectors as a function of time. Once we have them, we can recover the strain through the antenna pattern functions, which are obtained as described in Section 6.4:

$$\begin{aligned} F_+ &= (n \cdot u)(b \cdot u) - (n \cdot v)(b \cdot v) \\ F_\times &= (n \cdot u)(b \cdot v) + (n \cdot v)(b \cdot u). \end{aligned}$$

i Note 12: Scalar products in spherical coordinates

The scalar product between two unit vectors in spherical coordinates can be computed from their colatitudes and azimuths as follows:

$$a \cdot b = \cos \theta_a \cos \theta_b + \sin \theta_a \sin \theta_b \cos(\phi_a - \phi_b).$$

Computing these vectors can be somewhat expensive (on the scale of tens of milliseconds per evaluation), and their variation is relatively slow, on the scale of a month. Therefore, it is advantageous to pre-compute them on a grid, cache the results and interpolate them as required. We compute them in an ICRS frame with the lunarsky package [111]. We

¹As a sanity check, consider the relative signs of t_0 and this phase term: increasing t_0 means the introduction of a delay, and accordingly increasing the scalar product $\vec{r} \cdot \hat{n}$ means moving the detector further from the source, which also delays the signal’s arrival.




Figure 8.8.: Evolution of the response functions for the two LGWA channels over one month. The x and y channels correspond to two orthogonal horizontal directions across which displacement is measured.

store them in spherical coordinates, with which we can still directly compute any scalar products (Note 12).

i Note 13: The position of the LGWA

The precise location of the LGWA is a required input for the computation of these vectors as a function of time. It is defined by three angles: lunar latitude, lunar longitude and azimuth of one LGWA sensing direction, with the other being perpendicular to it.

Based on the results discussed in Section 8.1.1, we tentatively choose a lunar latitude of 88.5 degrees, a longitude of 290 degrees, and an azimuth of 0.

In order to evaluate how dense the interpolation grid must be to satisfy a given accuracy requirement, we compute it with different grid sizes for a duration of 10 years,² and for each grid size we:

- uniformly sample a moment in time t within the interpolation region;
- uniformly sample a unit vector u on the sphere;

²Specifically, from 2030 to 2040.

- compute the error in the scalar products, $|u \cdot (n^{\text{interp}}(t) - n^{\text{true}}(t))|$ and similarly for a and b .

As figure Figure 8.9 shows, we can achieve errors on the scale of 10^{-4} on the scalar products with only 10^4 points, which take about a minute to compute, and adding more as required is easy. Linear interpolation is sufficient to achieve this level of accuracy.




Figure 8.9.: Errors in scalar products computed with an interpolated version of the detector-centric basis vectors. The errors are computed from 400 randomly sampled points in the sky and time. Since the detector is located near the pole, the normal vector varies less across the month, and is therefore easier to interpolate.

In order to compute the strain in Equation 8.1 we also need the detector position. We apply the interpolation approach here as well, pre-computing the position on an equally spaced grid, and we can run a similar test to that in Figure 8.9 to determine how fine the interpolation grid needs to be.

The results in figure Figure 8.10 show that in order to achieve an error of 10^{-4} radians we need as many as 160 thousand interpolation points for the 10-year span we are considering. We compute the maximum phase error by comparing the position error to the angular wavelength corresponding to the smallest angular wavelength in the LGWA band (i.e. the highest frequency): $c/(2\pi \times 3\text{Hz}) \approx 1.6 \times 10^7\text{m}$. This is the worst-case scenario, as to achieve it we would need a source at the upper edge of the detector's band to also be perfectly aligned with the position error vector $\vec{r}_{\text{interp}} - \vec{r}_{\text{true}}$.




Figure 8.10.: Interpolation errors in position as a function of the number of interpolation points.

The position interpolation requires about 10 times more points to achieve a comparable level of accuracy, but each evaluation of the position of the detector is roughly 10 times faster than an evaluation of the basis vectors.

In Figure 8.11 we illustrate how the position vector encodes several “epicycles”: motion of the Earth-Moon system around the Sun, of the Moon around the Earth, and finally of the detector location from the lunar center.

8.1.4.1. The importance of computing time to merger correctly

In the Stationary Phase Approximation, the time left until merger for a source whose phase in the frequency domain is known is given by Equation 6.6.

One may think that, for the purpose of computing the time to merger, a precise description of the phase to a high PN order is not necessary, since high-PN order corrections only affect the late stages of the binary’s evolution. If we only consider the leading order in the phase (Equation 5.1), we get the following expression for the time to merger:

$$t_{\text{lowest order}}(f) = -\frac{5}{256\pi^{8/3}} \frac{1}{\mathcal{M}_c^{5/3}} \frac{1}{f^{8/3}}.$$




Figure 8.11.: Position of the detector as a function of time over 4 months. We show the evolution of the displacement vector from three different reference points. From the Solar System Barycenter we can only visually resolve a yearly modulation. From the Earth's center we can clearly see the effect of the lunar orbit. From the lunar center we see the displacement due to the detector's position vector. Due to the low tilt of the lunar rotation, this position vector is approximately constant. In all cases, the three Cartesian coordinates are computed in an ICRS frame.

The next-order correction to this expression for the phase comes with a factor of $v^2 \propto f^{2/3}$, which carries over during the differentiation, leading to the leading term in the error scaling with f^{-2} . In figure Figure 8.12 we show the error in time we get by only looking at the leading order: it can reach several hours to days in the deci-Hertz band for a light source.




Figure 8.12.: Error in the time to merger estimation by using the lowest-order Post-Newtonian term only. The source considered here is a binary with equal components, of 1.4 solar masses each, compatible with neutron stars such as those in GW170817.

The motion of the Moon over a day is comparable to the wavelength of a gravitational wave, band, hence incorrectly computing this time can lead to an error of several radians.

8.1.4.2. Projected waveforms

We end the discussion on waveform projection by analyzing the projected waveform for an injection modeled after GW250114. The pattern in amplitude is illustrated in Figure 8.13: the general scaling is $h_c \sim f^{-1/6}$, as expected since we are computing the characteristic strain amplitude of a source whose emission is dominated by low-order Post-Newtonian effects. On top of this trend we see modulations; their period in the time domain is a month (or rather half a month, since we are showing the absolute value of a sinusoid), and they are out of phase across the two LGWA channels, which measure displacement in two orthogonal horizontal direction.

Along with the injected waveform, we show the posterior predictive distribution obtained after performing an injection and recovery.




Figure 8.13.: Upper panel: Posterior Predictive Distribution and injected waveforms for our GW250114-like injection with the LGWA. The two colors represent the two horizontal measurement channels of the LGWA. We show the sensitivity curve of the LGWA in grey. Both are displayed as characteristic strains (Section 3.3.2). Lower panel: accumulation of SNR as a function of frequency, which can be referenced to time through the upper markers on the x axis. From Tissino et al. [183].

The phase of the response is more difficult to visualize, as it spans thousands of cycles during an observation. We illustrate one aspect of its variation by computing the ratio of two waveforms computed with identical parameters — again, the injection ones for GW250114 — but with a difference in 1 degree in right ascension. In the upper panel of Figure 8.14 we see the ratio of amplitudes for both channels, with a periodicity of one month. The phase, on the other hand, primarily exhibits variation with a slower period,

which corresponds to the yearly modulation driven by the lunar orbit around the Sun.³




Figure 8.14.: Ratio of projected waveforms, each compatible with GW250114, but with different right ascension by one degree.

8.2. Scientific prospects

8.2.1. Lunar science

The scientific potential of a very sensitive seismometer array at the lunar pole does not only include gravitational wave science [6]. Indeed, all the seismic motion which is a nuisance for gravitational wave observations would include valuable data on seismic events on the Moon. Here we give a brief glimpse into what problems may be probed thanks to LGWA data.

The seismometers deployed by the Apollo mission detected thousands of moonquakes, originating both near the surface (shallow moonquakes) and near the core (deep moonquakes). The information available comes from the lunar near side, but it seems to point to an asymmetry between it and the far side in terms of moonquake event rate. Whether this is a genuine effect or an artefact of selection bias is an open question.

³The phase is computed in a frame centered close to the Moon’s location at merger — for the details, see Section 8.3.

Furthermore, the LGWA would be able to precisely measure the excitation of lunar normal modes by Moonquakes.⁴ This would provide insights into the lunar interior structure, enabling the exploration of lunar origin models.

The formation mechanism of the Moon could also be investigated, as this is not a settled debate: the giant impact scenario, in which a large planetary object dubbed *Theia* impacted a young Earth forming a disk of debris which coalesced into our satellite, is the dominant theory, but it does have issues. The *isotopic crisis* is the problem posed by the observation that the Earth and Moon’s mantles have nearly identical ratios of isotopes such as $^{17}\text{O}/^{16}\text{O}$, $^{50}\text{Ti}/^{47}\text{Ti}$ and $^{182}\text{W}/^{184}\text{W}$. Simulations show that, under the giant impactor scenario, a majority of the Moon’s mass would be inherited from the impactor, while only a relatively small fraction of the Earth’s mass would be. Therefore, we expect different isotopic ratios: the Earth and the impactor, having formed independently, are likely to have different isotopic compositions. Indeed, the composition of other objects in the Solar System such as Mars is quite different from the Earth’s.

There are a number of theories proposed to solve this issue, including some mechanism for material equilibration between a proto-Earth and the orbiting disk that would become the Moon, and a different description of the impact, both in terms of improving numerical simulations and reevaluating their assumptions [122].

Mapping the composition of the lunar interior, which can be accomplished through seismic observations, will aid in the resolution of this problem.

8.2.2. Gravitational wave sources

8.2.2.1. Compact binaries

The range of gravitational wave sources observable in the most sensitive band for the LGWA, $[0.05, 2]\text{Hz}$, is quite wide. Compact binaries are a natural candidate: they are significant emitters of gravitational waves, and their study is a blossoming field due to the wealth of data from ground-based observations.

Figure 8.15 illustrates their time to merger as a function of mass. It is truncated at 10 years, an approximate upper bound for the LGWA mission’s duration. We events shown include: GW170817 and GW231123, respectively the lightest and heaviest out of the current population of *stellar-mass binaries* we detected. Depending on their mass, these may cross the deci-Hertz band hours to years before their merger.

⁴As discussed earlier, their excitation by gravitational waves would require an improbably strong signal.




Figure 8.15.: Time to merger as a function of mass, highlighting at which point various sources would cross the deciherzt band. For the purposes of the time to merger computation, the sources are assumed to have equal mass and no spin — including these effects changes the results, but not the overall qualitative picture. Furthermore, we only consider the dominant, $\ell|m| = 22$ mode of emission. We include the total mass of the three events considered in this chapter: GW170817, GW250114, and GW231123, with uncertainties as reported in the respective detection papers [4, 2, 1].




Figure 8.16.: The population of compact binaries in the GWTC-3 catalog, highlighting GW231123, compared to the noise curves of LGWA and ground-based detectors. From Iacovelli et al. [96].

In Figure 8.16 we further illustrate this aspect: a very massive binary like GW231123 would have been detectable by the LGWA, accumulating most of its SNR in the last day before merger.⁵

We can also see that the “area” between signal and noise⁶ is generally similar between LGWA and current ground based detectors. Indeed, a more systematic analysis confirms this. In Figure 8.17 we compute several realizations of a population compatible with the Gravitational Wave Transient Catalog, and compute their SNRs with various detectors, compared with the observed SNRs of the actual catalog. The SNRs obtained with the LGWA are indeed compatible, within catalog variance, with the original ones.

Despite their signal-to-noise ratios not being much higher than current observations, these low-frequency observations will provide qualitatively different information compared to high-frequency ones; we discuss this in more detail in Section 8.2.3.

The deci-Hertz band will also introduce the possibility to constrain the coalescence of

⁵This figure is meant to give an order of magnitude comparison between detector noise and signal magnitude, but since it does so with several different detectors and binaries it cannot be an apples-to-apples comparison: we should compare *projected* strain at the detector to the detector’s noise, but this would mean showing several different strain curves for every signal, as the projection onto each detector will be different. Here, instead, we are showing the amplitude of the strain polarizations before projection, $|h_+(f) + ih_\times(f)|$. Nevertheless, it communicates the correct qualitative message, as the variation in amplitude due to projection does not significantly change the signal magnitude.

⁶As we discussed in Section 3.3.2, the integral $\int (h_c/h_n)^2 d \log f$ gives us the square SNR; in this plot we show both signal and noise strains divided by $\sqrt{f}$, but this is immaterial as it leaves the integral unchanged.




Figure 8.17.: Cumulative SNR for the population of compact binaries in the GWTC catalogs. The detection rate for the LGWA at any given SNR is comparable to that obtained by the LVK. From Iacovelli et al. [96].

heavier binaries, which are inaccessible to ground-based observatories. Specifically, its sensitivity is optimal for binaries in the middle of the Intermediate-Mass Black Hole range, 100 to $10^5 M_\odot$. We illustrate this with a horizon plot in Figure 8.18: as a function of mass, we show the maximum redshift z (and corresponding luminosity distance d_L , under a fixed cosmology compatible with Planck data) to which an optimally oriented binary with that mass would be detectable.




Figure 8.18.: Horizon for equal mass black hole binaries as seen by the detectors for the next decade: Einstein Telescope, Lunar Gravitational Wave Antenna and LISA. The horizon is defined as the distance (or, equivalently, redshift) at which an optimally-oriented source would be detectable. From Ajith et al. [6].

The LGWA’s optimal sensitivity is for binaries with total detector-frame mass $(1 + z)M$ around $10^4 M_\odot$ (see Section 6.1.3). At the moment, full injections and recoveries for binaries in this range has not been performed. At these masses, the signals’ mergers will be in band, therefore (like with current observatories) several higher order modes will be necessary for a complete description of the signal. Stellar mass binaries are comparatively simpler, as their low-frequency emission is dominated by the $\ell|m| = 22$ mode, hence we chose them as the first target for our studies.

So far we discussed comparable-mass binaries, but Intermediate Mass Ratio Inspirals, binaries with mass ratios in the range $q = 10^{-2}$ to 10^{-4} , would also be a fascinating way to shed light into the formation of intermediate-mass black holes: the presence and orbits

of lighter black holes in the vicinity of an intermediate-mass black hole provide a precise probe of its environment, and in turn knowing about that can inform our understanding of its formation pathways.

Specifically, the LGWA’s operating frequency band allows for a long observation of these sources’ inspiral, which would allow for a precision measurement of any environmental effects, such as dynamical friction from a dark matter halo or migration torques from an accretion disk [41, 169]. Also, these early-time orbits would likely be non-circular, and measuring their eccentricity could provide evidence for or against different formation channels.

8.2.2.2. Non-CBC sources

Probing any previously unexplored region brings the possibility to find some completely new kinds of sources. These may have been previously theorized, their electromagnetic signatures may be known, but they may also be completely serendipitous new discoveries. While keeping in mind the last of these options is a possibility, it is useful to think of the candidate sources of gravitational waves in this band based on our current knowledge, as was done in the LGWA whitepaper [6] as well as in earlier work such as the review by Branchesi et al. [28].

Mergers of white dwarfs (WDs) are interesting events, as they may be the origin of certain supernovae. They generally happen when the binary is emitting gravitational waves in the deciHertz band. Their collision is not preceded by a measurable “chirp”, as the binary is still quite wide then, due to the large radius of white dwarfs: three to four orders of magnitude larger than black holes with the same mass. The nuclear detonation or deflagration which is expected as the two WDs impact each other is hard to model, with accurate 3D simulations of this scenario becoming available only recently [197, 199]. Collisions of WDs with other compact objects, such as neutron stars, are also of interest: they would emit at higher frequencies, which places them closer to the best sensitivity of LGWA around 0.3Hz, whereas double white dwarf systems generally collide before this point. The main limitation in the study of these systems is their low rate: they are comparatively rare in the local Universe, and the gravitational wave amplitudes they emit are not very high, meaning that they would need to be quite close — on the scale of a Megaparsec or less — to be detectable [20].

Tidal disruption events can also involve white dwarfs, for example a WD being disrupted after orbiting an intermediate mass black hole would produce a short-lived signal peaking in the deci-Hertz band [140, 184].

Core-collapse supernovae may also exhibit significant emission in the deci-Hertz band, but current simulations typically last for a few seconds, and are therefore unable to resolve this part of the spectrum [188].

Finally, the LGWA could contribute to the detection of a stochastic cosmological background of gravitational waves, but it would give significant results only by correlating the measurements of stations with a large separation, which would require an extension to the mission with seismometers placed at both lunar poles [Ajith et al. [6]; section 3.2.3].

8.2.3. Parameter estimation for stellar-mass binaries

We investigate the parameter estimation capabilities of the LGWA by performing injections and recoveries of some signals with parameters compatible with events detected by the LVK collaboration.

These injections are performed in zero noise, meaning that we investigate a Whittle likelihood ratio (Section 3.2) in the form

$$\log \mathcal{L}(\theta) = \Re(h(\theta)|h(\theta_0)) - \frac{1}{2}(h(\theta)|h(\theta)),$$

where the scalar products are computed according to the estimated LGWA strain noise PSD, which is computed based on the displacement noise PSD for the optimistic Silicon model in Figure 8.5 and the estimated response discussed in Section 8.1.3:

$$S_n^{\text{LGWA}}(f) = \frac{S_n^{\text{displacement}}(f)}{L(f)}.$$

We inject parameters compatible with two of the most well-known signals detected by the LVK collaboration:

- GW170817, the best-constrained neutron star binary detected with gravitational waves, and the only multi-messenger GW detection to date [4, 51];
- GW250114, the best-constrained black hole binary to date, whose detection by the LVK collaboration we discussed in detail in Chapter 4. We also provide comparisons with the injections with the Einstein Telescope discussed in Section 7.2.1.

We consider a 10-year observation of GW170817, which corresponds to the frequency band [208.1, 3000]mHz, and a 1-year observation of GW250114, in the band [27.1, 3000]mHz. For GW170817, the limiting factor is the mission lifetime. For GW250114 we could have chosen a longer observation, but as most of the SNR is accumulated in the last year before merger, this does not lead to significantly different results; truncating the observation this way highlights the fact that the constraints we see only require one year to be established.

The total SNR for GW170817 as seen by LGWA is 21.9, that of GW250114 is 34.8.

We perform these injections with the following software:

- the likelihood, including waveform generation and projection onto the LGWA, is handled within the [lgwa-response](#) package;
- prior distributions and interfacing across modules are handled through [bilby](#) [11];
- for sampling we use [nessai](#) [195].

Table 8.1.: Priors and injected values for our GW250114-like injection. The notation $[a, b] + c$ represents the interval $[a + c, b + c]$. The timing parameter Δt represents a deviation from the injected value, a GPS time for the merger of 1420878140.0 seconds, as measured at the location of the LGWA at this time. Phase is not included in this list as we do not explicitly sample over it, but by using a phase-marginalized likelihood we implicitly employ a uniform prior in $[0, 2\pi]$.

Parameter	Injected value	Unit	Prior function	Prior range
$\mathcal{M}$	31.2718	$M_\odot$	Uniform	$[-2, 2] \times 10^{-3} M_\odot + \mathcal{M}_0$
q	0.98		Uniform	$[0.125, 1]$
χ_1	-0.05		Aligned spin (Equation 5.2)	$[-0.9, 0.9]$
χ_2	0.01		Aligned spin (Equation 5.2)	$[-0.9, 0.9]$
d_L	414	Mpc	Uniform in the source frame (Equation 6.5)	$[100, 1000]$
θ_{JN}	0.7	rad	Sine	$[0, \pi]$
ψ	1.3	rad	Uniform	$[0, \pi]$
RA	2.3	rad	Uniform	$[0, 2\pi]$
DEC	0.2	rad	Cosine	$[-\pi/2, \pi/2]$
Δt	0	s	Uniform	$[-2, 2]$
m_1	36.3	$M_\odot$	Constraint	$[5, 100]$
m_2	35.5	$M_\odot$	Constraint	$[5, 100]$

For more details, see Tissino et al. [183] and the attached data release [181].

In Table 8.2 we outline the injected parameters for our GW170817-like binary neutron star injection. In this case, the posterior distribution occupies a significantly smaller volume; for computational efficiency, therefore, we perform these injections with restricted priors in several parameters, including the sky location. This is risky, since it may lead us to miss a mode in the posterior distribution on the sky location; while we expect no such additional

modes to exist due to the structure of the Doppler modulation in the observed waveform, this remains a potential systematic bias in our analysis. Future work will be required to relax this assumption.

Table 8.2.: Priors and injected values for our GW170817-like injection. The notation is the same as in Table 8.1, though for this injection we explicitly sample over phase.

Parameter	Injected value	Unit	Prior function	Prior range
$\mathcal{M}$	1.19757	$M_{\odot}$	Uniform	$[-2, 2] \times 10^{-6} M_{\odot} + \mathcal{M}_{\text{inj}}$
q	0.76		Uniform	[0.5, 1]
χ_1	0.02		Aligned spin (Equation 5.2)	[-0.8, 0.8]
χ_2	-0.01		Aligned spin (Equation 5.2)	[-0.8, 0.8]
Λ_1	414		Uniform	[0, 5000]
Λ_2	369		Uniform	[0, 5000]
d_L	37	Mpc	Uniform in the source frame (Equation 6.5)	[5, 100]
θ_{JN}	2.4	rad	Sine	[0, π]
ψ	0.4	rad	Uniform	[0, π]
RA	6.1	rad	Uniform	$[-0.01, 0.01]\text{rad} + \text{ra}_{\text{inj}}$
DEC	-0.6	rad	Cosine	$[-0.01, 0.01]\text{rad} + \text{dec}_{\text{inj}}$
ϕ	3.2		Uniform	[0, 2π]
$\Delta t_{\text{merger}}^{\text{center}}$	0	s	Uniform	[-3, 3]

8.2.3.1. Limitations

While they represent a significant improvement over Fisher-matrix estimates, our injections are not yet representative of a realistic data analysis pipeline for LGWA data, for several reasons.

The procedure we employ regarding response modelling is equivalent to performing inference assuming that

1. the lunar response is given by our model $L(f)$;

2. the lunar response is perfectly known during inference.

The first assumption is necessary — we need to establish a single reference model for the lunar response, though of course improvements to it will be useful, especially if they are driven by actual measurements from upcoming lunar missions. The second assumption can in principle be relaxed by injecting a signal computed with a response from our model, and recovering it with a variable response inferred simultaneously with the signal parameters, analogously to what is currently done for the calibration of ground-based detectors (Section 6.3).

Unfortunately, we have no reason to believe that a direct application of the smooth Gaussian process envelopes used for the inference of interferometric data (Section 6.3) would be appropriate in this context. On the contrary, a general feature observed in models and simulations of the lunar response is small-scale frequency dependence — the response is expected to be a “spiky” function of frequency, due to the combined effect of many local resonant modes. Defining an appropriate parametric model for it, therefore, is challenging, and it will require future work.

Another limitation of our analysis is a lack of data gaps. As we discussed in Section 8.1.2.1, moonquakes are likely to make sections of the data unusable, at an expected rate of roughly one hour every day. While this may not significantly reduce the overall SNR, it is not trivial computationally to perform a coherent analysis in this context. Burke et al. [37] thoroughly investigate the problem in the context of LISA, and it may be possible to investigate this problem using techniques borrowed from continuous wave science [176]. Nevertheless, in Section 8.3.2 we will present an argument as to why we do not expect our results in terms of source parameter constraints and sky localization to be significantly affected by this limitation.

Working directly with the zero-noise likelihood ratio simplifies the problem, meaning that we never have to leave the frequency domain; this is a blessing and a curse, since it also means that we are not performing a full pipeline test, which would need to start from simulated time-domain data. Furthermore, we are performing inference assuming that the signal’s presence is known, and no other signals nor non-Gaussian noise sources are present in the data. As we transition from forecasting to developing the analysis infrastructure for the mission, these issues can be addressed, eventually leading to a realistic simulation and blind analysis of simulated time-domain seismometer data.

Finally, the impact of working in the zero-noise case as opposed to generating a noise realization is comparatively minor. We explore this in more detail in Section 8.2.4.

8.2.3.2. Sky localization

Based on analytical results (Equation 2.17), we expect long observations to tightly constrain the sky position of the sources, due to the significant Doppler modulation of the signal throughout the year.

For GW170817 this is indeed the case, as shown in Figure 8.19: the posterior is extremely tight, and it remains so when repeating the same experiment with a lower upper limit in frequency of 710mHz, which corresponds to removing the last 14 days before merger. This only reduces the SNR from 21.9 to 21.3.




Figure 8.19.: Sky localization for a GW170817-like injection with the LGWA. Both for the full signal, and excluding the last two weeks of data, the area is tiny, on the order of a few tens of square arcminutes.

For GW250114, despite the higher SNR, the localization is much less precise. In Figure 8.20 we can see two nearby modes in the sky, corresponding respectively to a face-on source with $\theta_{JN} < \pi/2$ and a face-off source with $\theta_{JN} > \pi/2$. The LVK posterior distribution exhibits the same qualitative behavior, which in both cases is due to the sky position being also constrained by the relative amplitude of the signal in two detectors (for the LVK) and in the two LGWA channels at various times.




Figure 8.20.: Sky location, inclination angle and distance posteriors for a GW250114-like injection. We compare the results with those obtained with other detectors: the Einstein Telescope in its 2L configuration, and the posterior distribution provided by the LVK collaboration for the real event. From Tissino et al. [183].

The Einstein Telescope in its 2L configuration would measure this position with extreme precision, due to its SNR of more than 600. As we discussed in Section 7.2.1, ET in its Δ configuration would constrain the position of this source to two modes in the sky; if operational at the same time, the LGWA would easily allow us to discriminate among them.

8.2.3.3. Intrinsic parameter constraints

An observation of the gravitational waves emitted by a binary system constrains both of the masses, but as discussed in Section 5.1 the chirp mass is the lowest-order contribution, while the mass ratio has a higher PN order. For stellar mass binaries, the LGWA provides extremely precise constraints on the chirp mass, while its constraints on higher order parameters are not as tight.

In Figure 8.21 we show the posterior distribution on chirp mass, mass ratio and effective spin for GW250114. The LGWA measurement of chirp mass is extremely precise, with a 90% width of $0.0002M_{\odot}$; for reference, the interval for the Einstein Telescope is roughly 20 times wider, and the one in the original LVK analysis is 4000 times wider, at roughly $0.8M_{\odot}$. In contrast, the effective spin and mass ratio measurements from the LGWA have comparable precision with the LVK analysis, being respectively slightly narrower and wider. The Einstein Telescope, in contrast, exhibits extremely precise posterior distributions for both. The performance of the 2L and Δ configurations is almost identical for these parameters.

An important technical aspect to note is the fact that the relative precision on the detector-frame chirp mass of this observation, better than one part in 10^5 , requires us to precisely account for the Doppler shift between reference frames in a multi-detector context: as we discussed in Section 6.1.4, ground based observations often use a geocentric frame, which can be approximated as *moving with constant velocity*, but not as *stationary* compared to the SSB: at this precision level on the chirp mass, the distinction is appreciable.

While the chirp mass itself does not have a particular astrophysical significance, these analyses provide an indication of the LGWA’s sensitivity to tiny fluctuations at low PN order, which affect the long-duration frequency evolution of the signal. We expect this to mean it will be very effective at constraining other effects, such as contributions from dark matter in the environment, or eccentricity, both of which modify the phase of the binary at low frequencies.

8.2.4. Impact of non-zero Gaussian noise

Throughout this thesis we perform zero-noise injections, meaning that we assume the realization of the Gaussian noise process is the maximum-likelihood one, $n = 0$; however,




Figure 8.21.: Posterior distribution for the chirp mass, mass ratio and effective spin parameters of our GW250114-like injection. From Tissino et al. [183].

we still compute scalar products (Equation 3.3) accounting for the distribution it is drawn from, which we describe through the spectral density $S_n(f)$.

If the Fisher approximation holds, *i.e.* the posterior distribution is Gaussian in the parameters, then the effect of the noise realization is to bias the mean of the distribution away from its true value, but it will leave the covariance unchanged, as we can see from Equation 2.16.

Under realistic conditions, this will not be the case, and we expect to see noise-driven fluctuations in the covariance as well. However, we do expect the zero-noise posterior distribution $p(\theta|d = h_0)$ to be *representative* of the distributions $p(\theta|d = h_0 + n)$. In this section we investigate the impact of non-zero Gaussian noise realizations on our GW250114 injection with the Lunar Gravitational Wave Antenna.

Similar investigations have already been performed in other contexts, such as for example the ringdown of GW150914 [53] or the estimation of the neutron star equation of state [190]. Our results are qualitatively compatible with theirs.

We need to explore the space of distributions $p(\theta|d = h_0 + n)$ with varying noise $n \sim \mathcal{L}(n)$, which requires a representative number of noise realizations, at least in the hundreds.

Repeating the same injection we performed hundreds of times is not possible for us due to computational constraints. In order to circumvent this issue, we take two complementary approaches. The first (Section 8.2.4.1) is direct: we perform a campaign of 250 injections, all identical except for a different noise realization, and compare them to the corresponding zero-noise injection. To reduce the computational expense of the task, however, we perform these in a simplified scenario: we condition on the injected values for intrinsic parameters, use a phase-marginalized likelihood, and only sample over a reduced parameter space: time, sky position, distance, polarization and inclination angle, with priors as described in Table 8.1.

The second approach (Section 8.2.4.2) is more computationally efficient, and it allows us to explore the full parameter space, but it comes at the expense of a potential bias. We start from samples obtained in the zero-noise case, define the likelihood with a noise realization, and reweight them as described in Section 2.1.1.2, with weights

$$w_i = \frac{\mathcal{L}(d = h_0 + n|\theta_i)}{\mathcal{L}(d = h_0|\theta_i)}. \quad (8.2)$$

The concern, as usual with reweighting, is that this procedure might fail to accurately model regions in the parameter space which are not adequately covered by the noiseless posterior distribution. However, by reweighting the nested samples directly, we are able to obtain a satisfactory amount of effective samples.

In all cases, we are able to avoid simulationg full-length data by exploiting the properties of relative binning summary data, as described in Section 3.4.1.1.

The results obtained with the two approaches are discussed in Section 8.2.4.1 and Section 8.2.4.2 respectively.⁷

These injections are closer to a realistic scenario than the zero-noise case, but still not fully realistic: they are performed assuming that the noise is Gaussian, and colored according to a known power spectral density; neither of these will be true in practice.

8.2.4.1. Restricted noisy parameter estimation

We perform a series of 250 injections with the same prior assumptions and injected values reported in Table 8.1, but fixing chirp mass, mass ratio, and spins to their injected value. For each injection, we keep an identical analysis setup, and only vary the noise realization. Furthermore, we perform one extra reference analysis the same setup but with a noise realization $n = 0$.

In figure Figure 8.22 we show the distribution of summary statistics obtained in the noisy case, compared to the zero-noise one. As a general pattern, we can see that for most parameters the distribution of the noisy statistics is quite close to the zero-noise case. The width of the distribution varies significantly across parameters. For example, the polarization angle is nearly unconstrained in all cases, resulting in nearly identical intervals.

Other parameters exhibit more variation, but the value 1 (*i.e.* the summary statistic being equal to the zero-noise case) is generally quite close to the median of the distribution. The main outlier is the luminosity distance, for which the interval width in the noisy case is systematically larger than the zero-noise case: the difference is small (less than 10% additional width) but consistent across realizations. The reason for this effect is nonlinearity in the mapping between parameter and log-likelihood values; specifically, in this case, the degeneracy between distance and inclination.

As shown in Figure 8.23, the right boundaries of the distance 90% credible intervals in the noisy case are well-represented by the zero-noise case, while a sistematic difference arises in the left boundaries, which corresponds to a peak at low values of distance appearing in many of these distributions. This peak corresponds to a significant amount of posterior mass in the $\theta_{\text{JN}} \approx 0$ region, which is excluded in the zero-noise case.

⁷For clarity, we also note here that the results here are not expected to produce a diagonal p - p plot (Section 2.1.2): for that to be true, we would need to vary the injected value as well as the noise realization. For example, the mass ratio we are injecting is close to 1; with the typical uncertainties of this measurement, the injected value will always correspond to a high quantile of the distribution. This is expected and does not indicate a bias in our setup.




Figure 8.22.: Boxplot comparing several summary statistics between the noisy and zero-noise case, for our campaign of injections. $I_{90}(p)$ denotes the 90 percent symmetric interval for parameter p . $\Delta\Omega$ denotes the 90% HPD area in the sky. For all summary statistics, we report the ratio of their values in the noisy injections versus the zero-noise case. We show this as follows: an orange line denotes the distribution median, a box around it shows the 25-75th percentile range, whiskers show the 5th and 95th percentiles, and outliers beyond this range are individually shown as circles.




Figure 8.23.: Kernel density estimates of the marginal posterior distributions on luminosity distance in the zero-noise (thick, dashed) and noisy (thin, continuous) case. We highlight the boundaries of the symmetric 90% interval for each of these as vertical blue lines.

Finally, we investigate whether the values for the maximum likelihood, maximum posterior, evidence and posterior volume in the noisy case are consistent with theoretical expectations. It can be shown [88, appendix C] that the log-likelihood is expected to be distributed according to $\log \mathcal{L} \sim \mathcal{N}(\mu = \rho_{\text{opt}}^2/2, \sigma = \rho_{\text{opt}}^2)$, where ρ_{opt} is the optimal SNR (Equation 3.5). Then, when comparing with the noise-free case, we expect to see that *e.g.* the ratio of log-likelihood peaks will be distributed according to $\log \mathcal{L} / \log \mathcal{L}_{n=0} \sim \mathcal{N}(\mu = 1, \sigma = 2/\rho_{\text{opt}})$.




Figure 8.24.: Boxplot comparing sampling summary statistics obtained with our injection campaigns. For all summary statistics, we report the ratio of their values in the noisy injections versus the zero-noise case. The quantiles shown are defined as in Figure 8.22.

As shown in Figure 8.24, this expectation also holds when considering the maximum value of the log-posterior, as well as the log-evidence (for signal against noise, *i.e.* the integral of the residual likelihood). It appears that the tails of the distribution of these statistics in our noisy injections are slightly less populated than what one might expect based on a normal distribution.

8.2.4.2. Full noisy posterior reweighting

In order to estimate the effect of noise with the full parameter space, we reweight the samples obtained in zero noise for GW250114 (for injected values and priors see Table 8.1) with 1000 realizations of the noisy likelihood.

We can quantify the extent to the reweighted samples accurately describe the distribution they ought to be drawn from with the effective sample size (Equation 2.15). We start from 2.6×10^4 posterior samples, and as Figure 8.25 shows, we are sometimes left with less

than a hundred effective samples after the reweighting procedure: too few to make robust claims.

However, as that figure also shows, there is an alternative: we can perform the reweighting starting from the full set of 1.7×10^5 nested samples. These are already weighted, with weights given by Equation 2.9, which we can update by multiplying them by the factor in Equation 8.2.




Figure 8.25.: Sampling efficiencies for the reweighting approach, when starting from either equally-weighted posterior samples or nested samples.

In Figure 8.25 we compare the sample efficiencies obtained when starting from the nested samples to the ones we get with the posterior samples. The increase in median sample efficiency is by a factor 3 compared to what we get by weighting posterior samples, and the tail toward lower sample efficiencies is significantly less pronounced. Heuristically, the nested samples do not only cover the typical posterior volume in the baseline, zero-noise scenario, but also provide coverage of regions which are unlikely with it, but which may be more likely with a given noise realization.

In Figure 8.26 we show the fluctuations of some summary statistics describing the posterior distributions, analogously to Figure 8.22. The results are generally similar: the widths and sky areas in the zero-noise case are near the bulk of their distribution in the noisy case. The distance distributions are systematically wider in the noisy case than the zero-noise one, as observed in Figure 8.22.

For the sky area and coordinates, we observe a trend reversal: while in Figure 8.22 we typically observed narrower distributions in the noisy case than the zero-noise one, in Figure 8.26 we find them to be systematically broader. The cause of this is unclear: for the sky areas, it may be related to imprecision of the kernel density estimate we use when




Figure 8.26.: Boxplot comparing several summary statistics obtained with the reweighting approach for 1000 different noise realizations. For all summary statistics, we report the ratio of their values in the noisy injections versus the zero-noise case. The quantiles shown are defined as in Figure 8.22.

working with weighted samples. However, the trend is also observed in right ascension and declination alone: no density estimate is required to compute their quantiles, so it cannot be the sole cause.

Overall, when looking at the results obtained with our two approaches, we can say that the inclusion of Gaussian noise can modify posterior distributions in nontrivial ways, both by making them narrower and broader; which of these is more often the case depends on the specific scenario. Nevertheless, at least for this injection, the impact of the inclusion of noise in the injection is generally quite minor, inducing a scatter 90% quantile width on the order of a few tens of percent points at most.

8.3. Reference frames for long observations

In Section 6.1.4 we introduced the problem of choosing a reference frame to perform the analysis in. LGWA observations of stellar mass binaries are generally at least a month long, therefore the curvature of the lunar orbit around the Sun is non-negligible: a frame centered on the Earth or Moon would be measurably accelerated compared to the source.

This would seem to indicate that we should analyze our signals from the Solar System Barycenter frame; however, that is not actually necessary. We need a frame that is *comoving* with the SSB, but not necessarily *centered* there, and this is an important distinction for the LGWA. Let us consider again the example from Figure 8.14, only now we shall compare the phases in two different reference frames: one centered at the SSB, and one comoving with the SSB, but centered close to the Moon's location when it receive the merger. For clarity, we also switch to a linear scale in frequency.

In Figure 8.27 we can see how the phase dramatically changes, even for a small fluctuation in the parameters, if the center is at the SSB.

Furthermore, we can see that the change of phase is roughly linear in frequency. Such a phase shift is equivalent to a global time shift: consider a frequency domain signal whose phase includes a linear component, $\tilde{h}(f) = A(f)e^{i(\phi(f)+xf)}$ for some constant x . Its inverse Fourier transform will be

$$h(t) = \int df A(f) e^{i(\phi(f)+xf)-2\pi i f t} = \int df A(f) e^{i\phi(f)-2\pi i f (t-\delta t)},$$

where $\delta t = x/2\pi$. This is the Fourier transform of the signal without the linear term, $A(f)e^{i\phi(f)}$, evaluated at a time $t - \delta t$.

We can visualize what is happening: if the reference frame center is at the SSB, that is where all times are measured. In this example, we are moving the source position alone, keeping all other parameters constant, including the time. A signal which reaches the Sun




Figure 8.27.: Phase difference of projected waveforms, each compatible with GW250114, but with different right ascension by one degree, shown in two different frames.

at a time t_0 will reach our detector at a time $t_0 + \vec{r} \cdot \vec{k}/c$, as described in Section 6.1.4. Since $\vec{r}$ is large in magnitude, small perturbations in $\vec{k}$ change the arrival time of the signal significantly, which is what we are seeing here.

In practice, using the SSB as a frame center entails considerable degeneracies between the sky position and the reference time, which is also less precisely measured. The situation can be significantly improved by finding a position where time is well measured, so that this effect will be minimized. Fortunately, we can do so without repeating the analysis several times, as detailed in Section 8.3.1.

We combine this with another possible optimization, which regards the choice of reference event within our waveform. Typically, for the analysis of signals observed by ground-based interferometers, this event is the merger; however, this event is out of band for the LGWA, so one may wonder if there is a better choice. We can compute the SPA time at a given frequency based on Equation 6.6:

$$t_{\text{merger}} - t_{22}(f; \theta) = -\frac{1}{2\pi} \frac{d\phi_{22}(f; \theta)}{df}.$$

Then, we can parameterize our waveforms based on $t_{22}(f_0; \theta)$ as opposed to the merger time.




Figure 8.28.: An illustration of what it means to parameterize waveforms based on an evolutionary stage other than the merger. For each value of the chirp mass in range between 10 and 40 solar masses, we show the evolutionary track of frequency as a function of time. In the top panel, all curves correspond to binaries which reach coalescence at time 0 seconds, while in the bottom panel, they reach a frequency of 30Hz at time 0 seconds. These descriptions are equivalent, but the corresponding timing parameter may have a different variance, depending on the detector’s noise curve.

8.3.1. Shifting reference frame in post-processing

We start from the assumption of having successfully performed a parameter estimation run, in which we measured merger time referring to some center r_0 . We can convert each timing posterior sample to one measuring $t_{22}(f)$ at a different position r by the expression

$$t_i^{f,r} = t_i^{\text{mrg},r_0} - \frac{(r - r_0) \cdot \hat{n}_i}{c} + \frac{1}{2\pi} \frac{d\varphi_{22}(f; \theta_i)}{df},$$

where $\hat{n}$ is a unit vector pointing toward the source.

This gives us a new set of samples for the time, whose variance may be higher or lower than the baseline depending on the specifics of the signal. The timing uncertainty is a function of four variables $(\vec{r}, f)$, which we can visualize in a few different ways, using the injection with GW250114 described in Section 8.2.3 as our reference example.

Table 8.3.: Timing uncertainty for some choices of reference frame and evolutionary stage, for our GW250114-like injection. By “Earth” and “Moon” we refer to their location at a given time, our reference frames are always comoving with the Solar System Barycenter.

Location	Stage	Timing uncertainty (90% interval)
global optimum	$f_{22} = 0.56$ Hz	0.1 s
Earth	$f_{22} = 0.56$ Hz	0.15 s
Moon	merger	0.24 s
Earth	merger	0.25 s
Solar System Barycenter	merger	11.66 s
Solar System Barycenter	$f_{22} = 0.56$ Hz	11.72 s
Earth	$f_{22} = 0.03$ Hz	238 s

In Table 8.3 we show the timing uncertainty obtained at some locations and evolutionary stages. The global optimum in uncertainty is achieved at a location close to the position of the Moon a few hours from merger, as we will show in more detail in Figure 8.30. Qualitatively, we see that for reference frequencies at or above the best detector sensitivity (≈ 0.4 Hz), and locations close to the Earth-Moon system, the timing uncertainty remains low. On the contrary, at distant locations such as the Solar System Barycenter, or at very low reference frequencies, the timing uncertainty becomes very large.

In Figure 8.29 we vary the frequency f and, for each of its values, we compute the position r such that the 90% width of the samples $t^{f,r}$ is minimized. We then show these 90%




Figure 8.29.: Timing uncertainty by frequency for our GW250114-like injection. For all panels, the meaning of the vertical axis is above the panel itself. In the top panel we show the characteristic noise and signal strains as a function of frequency, as a reference. In the second panel from the top, we show the uncertainty in time as measured at a given reference frequency, as a function of that reference frequency. In the third panel we perform an analogous operation, but looking at the uncertainty in frequency as measured at a given reference time. From Tissino et al. [183].




Figure 8.30.: Spatial variation of the timing uncertainty for our GW250114-like injection. The geometry of the spatial slices we employ is described in the text. The projected trajectory and position of the Earth and Moon are represented with green and grey lines and circles. The projected optimal localization position for the row's frequency is shown as a red line. The projection of the Adapted from Tissino et al. [183].

widths as a function of frequency. One can see a minimum around 0.56Hz, just above the point where the LGWA's sensitivity is best.

In Figure 8.30, instead, we show the spatial variation of the timing uncertainty by fixing three values of f , which include the frequency of the global minimum, one lower frequency, and the merger. We also fix three orthogonal spatial slices, each passing through the global minimum, and compute the timing uncertainty in these slices, in order to show its 3D structure. We align the z axis in this frame with the injected source direction. The frequency-dependent minimum in Figure 8.29 is a curve in this visualization: $\vec{r}_{\min}(f)$. We show its projection onto each of the three planes in red, and similarly the trajectories of the Earth and Moon in green and grey. Finally, black crosses denote the projected location of the minimum uncertainty at the given frequency.

The shape of the iso-uncertainty ellipses generally extends further along the z axis than along the others. We give an analytical justification for this in Note 14.

i Note 14: Analytic calculation of the uncertainty spatial distribution

We start by re-parameterizing the sky coordinates, such that any sky position vector is given by $\hat{n} = (1 - \gamma)\hat{m} + \alpha\hat{u} + \beta\hat{v}$, where $\hat{m}$ is the unit vector in the direction of the posterior distribution average, while $\hat{u}$ and $\hat{v}$ complete an orthonormal basis with it, and are rotated so that for our posterior distribution the covariance between α and β is zero.

By this definition, α and β have zero mean. Then, we shift our time parameter t so that it also has zero mean. The parameter γ is always determined by α and β because $\hat{n}$ must have unit norm, and for a tight localization $|\gamma| \ll |\alpha|, |\beta|$, so γ -dependent terms are expected to be subdominant.

Nevertheless, we will include γ in our calculations.

Let us write down the covariance matrix for t , α and β . We assume to be able to find a location where it is diagonal:

$$\Sigma = \begin{pmatrix} \sigma_t^2 & 0 & 0 \\ 0 & \sigma_\alpha^2 & 0 \\ 0 & 0 & \sigma_\beta^2 \end{pmatrix}.$$

Now we will determine what happens to the timing variance under the transformation

$$\tilde{t} = t + \frac{\Delta\vec{r}}{c} \cdot (\hat{m}(1 - \gamma) + \alpha\hat{u} + \beta\hat{v}).$$

The new timing mean will be $\mathbb{E}[\tilde{t}] = \Delta\vec{r} \cdot \hat{m}(1 - \bar{\gamma})/c$, where $\bar{\gamma} = \mathbb{E}[\gamma]$, while the variance will be

$$\begin{aligned}
\mathbb{E} \left[\left(\tilde{t} - \frac{\Delta \vec{r} \cdot \hat{m}}{c} (1 - \bar{\gamma}) \right)^2 \right] &= \mathbb{E} \left[\left(t + \frac{\Delta \vec{r}}{c} \cdot (\hat{m}(\bar{\gamma} - \gamma) + \alpha \hat{u} + \beta \hat{v}) \right)^2 \right] \\
&= \sigma_t^2 + \left(\frac{\Delta \vec{r} \cdot \hat{u}}{c} \right)^2 \sigma_\alpha^2 + \left(\frac{\Delta \vec{r} \cdot \hat{v}}{c} \right)^2 \sigma_\beta^2 + \\
&\quad + \left(\frac{\Delta \vec{r} \cdot \hat{m}}{c} \right)^2 \mathbb{E} [(\bar{\gamma} - \gamma)^2] \\
&\quad - 2 \frac{\Delta \vec{r} \cdot \hat{m}}{c} \left(\mathbb{E} [t(\gamma - \bar{\gamma})] \right. \\
&\quad \left. + \frac{\Delta \vec{r} \cdot \hat{u}}{c} \mathbb{E} [(\gamma - \bar{\gamma})\alpha] + \frac{\Delta \vec{r} \cdot \hat{v}}{c} \mathbb{E} [(\gamma - \bar{\gamma})\beta] \right).
\end{aligned}$$

Neglecting the terms depending on γ , which is subdominant for a relatively tight localization, we can recognize the ellipses from Figure 8.30 if we constrain ourselves to the $(\hat{u}, \hat{v})$ plane.

Finally, in Figure 8.31 we show the velocity of the point $\vec{r}_{\min}(t(f))$, where we compute the time $t(f)$ associated with the frequency f based on a SPA of the injected signal. We can see that, even though its motion is obtained exclusively by post-processing posterior samples, it retains a certain “memory” of the detector’s trajectory through the Solar System.

8.3.2. Spanned area scaling validation

Fisher matrix results (Equation 2.17) suggest that the sky localization obtained with a long observation — all else being equal — scales with the inverse of the “area spanned by the detector’s motion”, projected onto the plane orthogonal to the source’s propagation direction.

As discussed in the seminal paper of Wen and Chen [192], this area should be computed in terms of the weighted variance and covariance of the detector motion, as outlined in Equation 2.18. We label this the *statistical area* A_s .

If the accumulation of SNR is uniform over time, this area is well-approximated by the convex hull of the projected trajectory, *i.e.* the set of points P in the plane for which one can find two times t_1 and t_2 such that the line connecting the projected positions $r(t_1)$ and $r(t_2)$ passes through P . We label the area of this set the *geometric area* A_g . The analytical expression in Equation 8.3 provides an accurate approximation of A_g , to the percent level.




Figure 8.31.: Velocity of the detector as a function of frequency, compared to the relative velocity of the detector and the minimum in timing uncertainty. For a long duration, the minimum approximately moves with the detector, as evidenced by their relative velocity being smaller in modulus than the Moon's orbital velocity.

$$A_a = \frac{(\omega_* t - \sin(\omega_* t)) |\hat{k} \cdot \hat{n}_{\text{ecl}}|}{2} \quad (8.3)$$

The geometric area has a direct and intuitive interpretation, while computing the statistical area is generally more involved. In this section, we validate the expectation that the sky localization scales inversely with the statistical area A_s , as well as exploring when the geometric area can be used to approximate A_s . In order to do so, we set up 12 injections with the same parameters as our baseline GW250114 analysis, except for the following:

1. The lower frequency bound is always kept at 27.1mHz;
2. the upper frequency bound is changed, so that the i -th injection lasts i months in band;
3. the luminosity distance d_L is adjusted so that the SNR for each injection is constant: this requires all but the 12 month injection to be unrealistically close.

The corresponding waveforms for the 11 and 12-month cases are shown in Figure 8.32.

i Note 15: The SNR integrand

We complement Figure 8.33 with an indication of how SNR is accumulated over time for each of the injections. Since our signal is composed of a single time-frequency track, we can write the integral which defines the total square SNR as

$$\text{SNR}^2 = \int \frac{|h(f)|^2 df}{\underbrace{S_n(f)}_{w(t)} dt}.$$

The average of $w(t)$, taken as a function of time, is equal to the total squared SNR. Figure 8.33 shows how $w(t)$ varies throughout the orbit for all our injections.

We perform these injections in three setups:

1. recovering **all parameters** except for spin;
2. recovering only sky position, time, phase (**simplified**);
3. modifying the LGWA noise curve so that each month accumulates an equal amount of SNR (**simplified, uniform SNR**).

The statistical area is quite close to the geometric one for the shorter injections; on the other hand, as we reach higher frequencies and a full year of observation, they diverge significantly. The reason for this can be visually seen in Figure 8.33, where we draw the areas, as well as showing the SNR accumulation over time along the trajectory, computed as described in Note 15.




Figure 8.32.: Waveforms for the 11 and 12 month normalized injections, shown in characteristic strain, together with the LGWA PSD and the artificial PSD used in the “uniform SNR” scenario. From Tissino et al. [183].

In the 12-month case, as Figure 8.13 also shows, most of the SNR is accumulated during the last week of observation, a comparatively small fraction of the orbit.

The shapes of the localization areas are shown in Figure 8.34, while Figure 8.35 summarizes the comparison between these setups and the Fisher matrix expression.

The bimodality in sky position which we observe in the full injection (Figure 8.20) is only retained in the 12 month case with all parameters. The simplified scenarios fix the inclination angle and are therefore unable by construction to reproduce it. In the runs lasting 11 months or less with all parameters, the SNR is more evenly distributed across the observation, enabling more precise localization through phase alone, thus removing the degeneracy with inclination.

As a general pattern, we see how the qualitative Fisher matrix scaling is validated for all but the 12 month case: the sky localization area decreases with the number of months, even though the SNR is kept constant, driven mostly by the change in area; the frequency increases slightly through the injections, but not nearly enough to explain the effect. The Fisher matrix approach systematically underestimates the localization areas; we attribute this to the fact that by construction it is only marginalized over arrival time, while even in our **simplified** scenario we also marginalize over phase. We did not simplify this sce-




Figure 8.33.: Top left panel: trajectories and convex hulls for the 12 month normalized injections, projected onto a plane orthogonal to the gravitational wave propagation direction, shown in the same color. We compare these to the statistical areas (Equation 2.18), shown as ellipses in pink. The line thickness for the trajectory is proportional to the SNR integrand as a function of time, defined in Note 15. Bottom left panel: SNR integrand as a function of time. Right panel: same as the top left, but zoomed in onto the 12 month injection. From Tissino et al. [183].




Figure 8.34.: Sky localization areas for the 12 month normalized injections in the three scenarios described in Section 8.3.2.

nario further, because conditioning over the arrival phase leads to pathological, “banded” multimodal posterior distributions in the sky.

The failure of the prediction in the 12 month case, with the localization worsening compared to the 11-month one, is due to the non-uniform accumulation of SNR, with most of it coming from the last few days of observation: while the geometric area A_g increases monotonically with time, the statistical area A_s decreases in this case. We validate this through the **simplified, uniform SNR** scenario, in which the SNR is accumulated at a constant rate over time — we obtain this by artificially setting a PSD in the form $S_n(f) \propto f^{4/3}$, as shown in Figure 8.32. This way, the SNR integrand $w(t)$ (Note 15) will approximately have the same distribution every month. In this scenario, we see that the localization sky area keeps decreasing even for the 12 month-long injection.

Finally, in Figure 8.35 we also see that the sky area for our 12-month **simplified, uniform SNR** injection, while smaller than the 11-month case, is still much larger than the Fisher matrix prediction using the geometric area A_g . This discrepancy is due to the failure of another assumption we are making: the monochromaticity one. We are using Equation 2.17 setting as the reference frequency the SNR-weighted average frequency, *i.e.*

$$\rho^2 f_0^2 \approx \int f^2 \frac{|h(f)|^2}{S_n(f)} df$$

In order to show a single curve in the plot, we are computing this with the LGWA PSD; in




Figure 8.35.: Fisher matrix expectation and measured sky localization, in the three scenarios described in Section 8.3.2. The **all parameters** and **simplified** scenarios use the correct LGWA PSD; in the former, we sample on all source parameters except for spin, in the latter, we only sample over sky position, arrival time and phase. The **simplified, uniform SNR** case is the same as the **simplified** one, except for the usage of an artificial PSD, constructed so that the SNR accumulated in each month is approximately equal. We compare these results to the Fisher matrix prediction, computed both using the statistical and (incorrectly) the geometric area.

the **simplified, uniform SNR** case this would give a lower result, and hence the estimated sky area would be higher. Specifically, in the 12-month case the estimated f_0 with this procedure is 391 mHz, while using the modified $f^{4/3}$ PSD we get a lower result by an order of magnitude, 33 mHz. Accounting for this, the scaling of the A_g Fisher prediction agrees with the **simplified, uniform SNR** results.




Figure 8.36.: Top panel: scaling of the geometric and statistical areas A_g and A_s as a function of the number of observation months, compared with the analytical approximation A_a . A_g and A_a agree in all cases, while A_s is generally smaller. Bottom panel: the ratio $\alpha(T) = A_s/A_g$ as a function of the number of observation months, on a logarithmic vertical scale. For both panels, we also compute the statistical area for a variety of total binary masses, leaving all other parameters unchanged; the total binary mass is shown with color, and it ranges from $M \approx 2M_\odot$ (a typical value for a neutron star binary) and $M \approx 72M_\odot$ (the value for our GW250114-like injection). This shows how the statistical area becomes a closer approximation to the geometric one as the total mass decreases. From Tissino et al. [183].

This analysis allows us to determine the crucial features of the signals which a deci-Hertz detector will precisely localize: they should be observed for a long time, at least a *few months*, and their SNR should be *well-distributed* across the observation. In the case of

CBCs, this points to light sources, with component masses $M \lesssim 30M_{\odot}$: we illustrate this in Figure 8.36, where we compute the statistical area for different total binary masses. Heavier binaries will generally have high SNRs, and their intrinsic parameters will be well-constrained, but their localization will not be very precise.

Another conclusion we can draw from this analysis is that the hour-long data gaps the LGWA is expected to experience (Section 8.1.2.1), though not yet included in our simulations (see Section 8.2.3.1), will not significantly impact sky localization, as long as we can maintain phase coherence across them. This is because the area spanned by the observation is not significantly affected if a small fraction of the orbit is removed.

8.3.3. Effect on sampling

In Section 8.3.1 we discussed how the reference frame can be changed *a posteriori* based on the results from an analysis. Indeed, any frame chosen is formally equivalent to any other, but this equivalence does not carry over to sampling efficiency: on the contrary, it is significantly more efficient to sample using a frame for which the timing uncertainty is low.

Demonstrating this requires us to perform at least one parameter estimation run with a suboptimal frame. As we shall shortly show, this entails a significant computational multiplier: doing so for the complete injection, which requires a good amount of resources even with a close-to-optimal frame, would be wasteful and lengthy. Therefore, we consider a simplified scenario, still based on the GW250114-like injection discussed in Section 8.2.3. We employ a phase-marginalized likelihood (Section 3.5.1) and only sample over sky position and time. We adopt a uniform prior in the sky, and a uniform prior in the timing parameter, whose width is *a posteriori* chosen to be approximately 6 times wider than the 90% credible interval, and centered at the true value, to ensure a fair comparison.

We compare two reference frame centers: one at the SSB, and one at the optimal location, as computed through a timing uncertainty minimization procedure. The required prior width on the timing parameter in these two cases is significantly different. Being able to use a narrower prior is part of the reason why sampling is faster in the optimal frame, but it is not the major contributor: to demonstrate this, we include an intermediate scenario where we sample based on the optimal frame, but with the same wide prior that would be required for the SSB frame.

Finally, since this comparison is about sampler performance, we compare two different samplers: dynesty [168] and nessai [195]. Both are commonly used in gravitational wave science. The proposals are obtained through a decomposition of the live points into ellipses with the former, and through a normalizing flow with the latter.

Table 8.4.: Variation in computational cost for different choices of reference frame center and sampler.

Likelihood evaluations	Sampling time [min]	Timing uncertainty [s]	Prior width [s]	Sampler	Center
91515	3	0.075	0.5	nessai	optimal
107878	3.6	0.075	27	nessai	optimal
659816	16.7	4.733	27	nessai	SSB
806067	16.8	0.074	0.5	dynesty	optimal
1530906	33	0.076	27	dynesty	optimal
9144813	198.6	4.766	27	dynesty	SSB

We show the results in Table 8.4. The difference in timing uncertainty between the two reference frames is about a factor 60, and the corresponding ratio of sampling times is about a factor 10 with dynesty, and a factor 5 with nessai. The intermediate scenario, using a wide prior but the optimal frame, is barely slower with nessai while it doubles the computational cost with dynesty.

Many variables can affect sampling time. To ensure that our results are robust, we also verify that the same scaling is observed in the number of likelihood evaluations.

In practice, the optimal reference frame center will be *a priori* unknown. The strategy we propose is to perform the first parameter estimation run with a heuristically-determined origin, compute the optimal origin, which should always be used thereafter. The heuristic we employed was to use the position of the detector when it receives the merger as a reference.

9. Conclusions

This thesis summarizes several novel findings in gravitational wave science, related to the detection and analysis of recent gravitational wave events and the preparation of data analysis techniques required for future detectors.

In January 2025 the LIGO-Virgo-KAGRA collaboration detected GW250114 [2], a merger of black holes which set a new record for clarity, as measured by its signal-to-noise ratio of nearly 80. Chapter 4 outlines how we used it to perform tests of General Relativity, constraining the frequency of the first overtone and validating Hawking’s area theorem. Working with such a precisely-measured signal also required a reexamination of the analysis assumptions that were previously made, which led to the identification and correction of a long-standing issue related to data normalization (Section 3.2.1).

Data analysis with the next generation of ground based detectors, such as the Einstein Telescope, will require efficient parameter estimation techniques. In Chapter 7, we show that parameter estimation on a population scale is possible thanks to neural posterior estimation (Section 2.3.2). While this work only included a subset of the sources we expect to be present in Einstein Telescope data, namely massive black hole binaries, it paves the way for the fast analysis of all the compact binary signals detected. Furthermore, it enables the comparison of different designs for the detector: we investigate the parameter constraints the Einstein Telescope would provide in several of its possible configurations. Our methods are fully Bayesian, and they allow us to illustrate how simple summary statistics like the variance in parameter estimates fail to fully capture the uncertainty in the inferred location of sources. For example, we show how the posterior distributions obtained with a triangular detector generally exhibit 8-fold multimodality in the sky for the sources in our sample, while a 2L configuration often yields multimodal estimates of the source’s distance.

The Lunar Gravitational Wave Antenna (LGWA, Chapter 8) will measure gravitational waves with an array of seismometers, a different principle compared to laser interferometers, in which the Moon itself is used as a resonant bar. Thanks to the quiet environment in permanently shadowed regions at the lunar pole, our seismometers will be able to make precise measurements at lower frequencies. We explore the science case of this detector, which ranges from lunar science to massive compact binaries of intermediate-mass black holes, light compact binaries such as those containing white dwarfs or neutron stars, and

stellar-mass black hole binaries. The LGWA will probe the early inspiral of binaries in the population currently being observed by the LIGO-Virgo-KAGRA detectors, providing early warning and precise constraints on their evolution and sky position. The tight constraints on the chirp mass indicate that other, more astrophysically relevant parameters such as eccentricity or environmental matter density would also be tightly constrained, though that remains a topic for future work. We investigate the way the LGWA localizes its sources, and find it to be primarily through the measurement of Doppler phasing as the detector orbits the Sun, which is more effective for less massive sources.

We show that straightforwardly employing standard data analysis techniques to these binaries can lead to significant inefficiency: through a minimization procedure, requiring little computational overhead, we are able to find an optimized coordinate frame, which reduces sampling time by an order of magnitude. This procedure is of general applicability, and it will allow for coherent multi-band data analysis.

References

- [1] A. G. Abac et al. “GW231123: A Binary Black Hole Merger with Total Mass 190-265 Msun”. In: *The Astrophysical Journal Letters* 993.1 (Oct. 2025), p. L25. ISSN: 2041-8205. DOI: [10.3847/2041-8213/ae0c9c](https://doi.org/10.3847/2041-8213/ae0c9c). URL: <https://doi.org/10.3847/2041-8213/ae0c9c> (visited on 02/19/2026).
- [2] A. G. Abac et al. “GW250114: Testing Hawking’s Area Law and the Kerr Nature of Black Holes”. In: *Physical Review Letters* 135.11 (Sept. 10, 2025), p. 111403. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/kw5g-d732](https://link.aps.org/doi/10.1103/kw5g-d732). URL: <https://link.aps.org/doi/10.1103/kw5g-d732> (visited on 10/06/2025).
- [3] Adrian Abac et al. “The Science of the Einstein Telescope”. In: *Journal of Cosmology and Astroparticle Physics* 2026.03 (Mar. 2026), p. 081. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2026/03/081](https://doi.org/10.1088/1475-7516/2026/03/081). URL: <https://doi.org/10.1088/1475-7516/2026/03/081> (visited on 04/25/2026).
- [4] B. P. Abbott et al. “GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral”. In: *Physical Review Letters* 119.16 (Oct. 16, 2017), p. 161101. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.119.161101](https://link.aps.org/doi/10.1103/PhysRevLett.119.161101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.119.161101> (visited on 04/04/2021).
- [5] P. Ajith et al. “Data Formats for Numerical Relativity Waves”. In: *ArXiv07090093 Gr-Qc* (Feb. 24, 2011). arXiv: [0709.0093 \[gr-qc\]](http://arxiv.org/abs/0709.0093). URL: <http://arxiv.org/abs/0709.0093> (visited on 03/03/2021).
- [6] Parameswaran Ajith et al. “The Lunar Gravitational-wave Antenna: Mission Studies and Science Case”. In: *Journal of Cosmology and Astroparticle Physics* 2025.01 (Jan. 2025), p. 108. ISSN: 1475-7516. DOI: [10.1088/1475-7516/2025/01/108](https://dx.doi.org/10.1088/1475-7516/2025/01/108). URL: <https://dx.doi.org/10.1088/1475-7516/2025/01/108> (visited on 07/30/2025).
- [7] Simone Albanesi et al. *Effective-One-Body Modeling for Generic Compact Binaries with Arbitrary Orbits*. Mar. 18, 2025. DOI: [10.48550/arXiv.2503.14580](https://arxiv.org/abs/2503.14580). arXiv: [2503.14580 \[gr-qc\]](https://arxiv.org/abs/2503.14580). URL: <http://arxiv.org/abs/2503.14580> (visited on 03/14/2026). Pre-published.
- [8] M. Andrés-Carcasona and G. Caneva Santoro. *No Love for Black Holes: Tightest Constraints on Tidal Love Numbers of Black Holes from GW250114*. Dec. 1, 2025. DOI: [10.48550/arXiv.2512.01918](https://arxiv.org/abs/2512.01918). arXiv: [2512.01918 \[gr-qc\]](https://arxiv.org/abs/2512.01918). URL: <http://arxiv.org/abs/2512.01918> (visited on 03/09/2026). Pre-published.

- [9] Greg Ashton et al. “Nested Sampling for Physical Scientists”. In: *Nature Reviews Methods Primers* 2.1 (May 26, 2022), pp. 1–22. ISSN: 2662-8449. DOI: [10.1038/s43586-022-00121-x](https://doi.org/10.1038/s43586-022-00121-x). URL: <https://www.nature.com/articles/s43586-022-00121-x> (visited on 11/23/2024).
- [10] Gregory Ashton. *Reconstructing and Resampling: A Guide to Utilising Posterior Samples from Gravitational Wave Observations*. Oct. 13, 2025. DOI: [10.48550/arXiv.2510.11197](https://doi.org/10.48550/arXiv.2510.11197). arXiv: [2510.11197](https://arxiv.org/abs/2510.11197) [gr-qc]. URL: <http://arxiv.org/abs/2510.11197> (visited on 10/20/2025). Pre-published.
- [11] Gregory Ashton et al. “Bilby: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy”. In: *The Astrophysical Journal Supplement Series* 241.2 (Apr. 2019), p. 27. ISSN: 0067-0049. DOI: [10.3847/1538-4365/ab06fc](https://doi.org/10.3847/1538-4365/ab06fc). URL: <https://doi.org/10.3847/1538-4365/ab06fc> (visited on 03/18/2026).
- [12] Augustine Kong. *A Note on Importance Sampling Using Standardized Weights*. Tech. Rep 348. University of Chicago, Dept. of Statistics, 1992. URL: <https://d3qi0qp55mx5f5.cloudfront.net/stat/docs/tech-rpts/tr348.pdf> (visited on 03/03/2026).
- [13] S. Babak et al. “Searching for Gravitational Waves from Binary Coalescence”. In: *Physical Review D* 87.2 (Jan. 22, 2013), p. 024033. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.87.024033](https://doi.org/10.1103/PhysRevD.87.024033). arXiv: [1208.3491](https://arxiv.org/abs/1208.3491) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/1208.3491> (visited on 05/02/2023).
- [14] Emily Baird et al. “Degeneracy between Mass and Spin in Black-Hole-Binary Waveforms”. In: *Physical Review D* 87.2 (Jan. 22, 2013), p. 024035. DOI: [10.1103/PhysRevD.87.024035](https://doi.org/10.1103/PhysRevD.87.024035). URL: <https://link.aps.org/doi/10.1103/PhysRevD.87.024035> (visited on 03/18/2026).
- [15] A. Makai Baker et al. “Significant Challenges for Astrophysical Inference with Next-Generation Gravitational-Wave Observatories”. In: *Phys. Rev. D* 112.10 (2025), p. 102004. DOI: [10.1103/n6t6-5wn3](https://doi.org/10.1103/n6t6-5wn3).
- [16] Biswajit Banerjee et al. “Camelidae on BOAT: Observation of a Second Spectral Component in GRB 221009A”. In: *arXiv.org* (May 24, 2024). URL: <https://arxiv.org/abs/2405.15855v1> (visited on 10/09/2024).
- [17] Jacob D. Bekenstein. “Generalized Second Law of Thermodynamics in Black-Hole Physics”. In: *Physical Review D* 9.12 (June 15, 1974), pp. 3292–3300. DOI: [10.1103/PhysRevD.9.3292](https://doi.org/10.1103/PhysRevD.9.3292). URL: <https://link.aps.org/doi/10.1103/PhysRevD.9.3292> (visited on 03/13/2026).
- [18] Enis Belgacem, Michele Maggiore, and Thomas Moreau. *Coupling Elastic Media to Gravitational Waves: An Effective Field Theory Approach*. Mar. 25, 2024. arXiv: [2403.16550](https://arxiv.org/abs/2403.16550) [astro-ph, physics:gr-qc, physics:hep-th]. URL: <http://arxiv.org/abs/2403.16550> (visited on 03/27/2024). Pre-published.

- [19] A. Ben-Menahem. “Excitation of the Earth’s Eigenvibrations by Gravitational Radiation from Astrophysical Sources”. In: *Il Nuovo Cimento C* 6.1 (Jan. 1, 1983), pp. 49–71. ISSN: 0390-5551. DOI: [10.1007/BF02511372](https://doi.org/10.1007/BF02511372). URL: <https://doi.org/10.1007/BF02511372> (visited on 03/14/2025).
- [20] Giovanni Benetti et al. “Observing Double White Dwarfs with the Lunar GW Antenna”. In: (Sept. 2025).
- [21] Emanuele Berti et al. *Black Hole Spectroscopy: From Theory to Experiment*. May 29, 2025. DOI: [10.48550/arXiv.2505.23895](https://arxiv.org/abs/2505.23895). arXiv: [2505.23895 \[gr-qc\]](https://arxiv.org/abs/2505.23895). URL: <http://arxiv.org/abs/2505.23895> (visited on 06/10/2025). Pre-published.
- [22] Alessandro Bertolini et al. “Mechanical Design of a Single-Axis Monolithic Accelerometer for Advanced Seismic Attenuation Systems”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 556.2 (Jan. 2006), pp. 616–623. ISSN: 01689002. DOI: [10.1016/j.nima.2005.10.117](https://doi.org/10.1016/j.nima.2005.10.117). URL: <https://linkinghub.elsevier.com/retrieve/pii/S0168900205020930> (visited on 04/11/2023).
- [23] Xiaoming Bi and Jan Harms. “Response of the Moon to Gravitational Waves”. In: *Physical Review D* 110.6 (Sept. 9, 2024), p. 064025. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.110.064025](https://doi.org/10.1103/PhysRevD.110.064025). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.064025> (visited on 03/14/2025).
- [24] *Bilby.Gw.Detector.Interferometer.Interferometer — Bilby 2.8 Documentation*. URL: https://bilby-dev.github.io/bilby/api/bilby.gw.detector.interferometer.Interferometer.html#bilby.gw.detector.interferometer.Interferometer.whiten_frequency_series (visited on 03/12/2026).
- [25] Luc Blanchet. “Post-Newtonian Theory for Gravitational Waves”. In: *Living Reviews in Relativity* 27.1 (July 10, 2024), p. 4. ISSN: 1433-8351. DOI: [10.1007/s41114-024-00050-z](https://doi.org/10.1007/s41114-024-00050-z). URL: <https://doi.org/10.1007/s41114-024-00050-z> (visited on 03/08/2026).
- [26] Alice Bonino et al. *Inferring Eccentricity Evolution from Observations of Coalescing Binary Black Holes*. July 21, 2022. arXiv: [2207.10474 \[gr-qc\]](https://arxiv.org/abs/2207.10474). URL: <http://arxiv.org/abs/2207.10474> (visited on 10/03/2022). Pre-published.
- [27] Michael Boyle, Robert Owen, and Harald P. Pfeiffer. “Geometric Approach to the Precession of Compact Binaries”. In: *Physical Review D* 84.12 (Dec. 2, 2011), p. 124011. DOI: [10.1103/PhysRevD.84.124011](https://doi.org/10.1103/PhysRevD.84.124011). URL: <https://link.aps.org/doi/10.1103/PhysRevD.84.124011> (visited on 03/18/2026).
- [28] Marica Branchesi et al. “Lunar Gravitational-Wave Detection”. In: *Space Science Reviews* 219.8 (Oct. 20, 2023), p. 67. ISSN: 1572-9672. DOI: [10.1007/s11214-023-01015-4](https://doi.org/10.1007/s11214-023-01015-4). URL: <https://doi.org/10.1007/s11214-023-01015-4> (visited on 10/20/2023).

- [29] Marica Branchesi et al. *Science with the Einstein Telescope: A Comparison of Different Designs*. Mar. 28, 2023. DOI: [10.48550/arXiv.2303.15923](https://doi.org/10.48550/arXiv.2303.15923). arXiv: [2303.15923](https://arxiv.org/abs/2303.15923) [gr-qc]. URL: <http://arxiv.org/abs/2303.15923> (visited on 03/29/2023). Pre-published.
- [30] Matteo Breschi et al. *Kilohertz Gravitational Waves from Binary Neutron Star Mergers: Inference of Postmerger Signals with the Einstein Telescope*. May 20, 2022. arXiv: [2205.09979](https://arxiv.org/abs/2205.09979) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2205.09979> (visited on 10/03/2022). Pre-published.
- [31] Matteo Breschi et al. “Kilohertz Gravitational Waves from Binary Neutron Star Mergers: Numerical-relativity Informed Postmerger Model”. In: *Phys. Rev. D* 109.6 (2024), p. 064009. DOI: [10.1103/PhysRevD.109.064009](https://doi.org/10.1103/PhysRevD.109.064009).
- [32] Matteo Breschi et al. “kiloHertz Gravitational Waves from Binary Neutron Star Remnants: Time-Domain Model and Constraints on Extreme Matter”. In: *Phys. Rev. D* 100.10 (2019), p. 104029. DOI: [10.1103/PhysRevD.100.104029](https://doi.org/10.1103/PhysRevD.100.104029).
- [33] R. Broucke and P. Cefola. “A Note on the Relations between True and Eccentric Anomalies in the Two-Body Problem”. In: *Celestial Mechanics* 7 (Apr. 1, 1973), pp. 388–389. ISSN: 0008-8714. DOI: [10.1007/BF01227859](https://doi.org/10.1007/BF01227859). URL: <https://ui.adsabs.harvard.edu/abs/1973CeMec...7..388B> (visited on 03/13/2026).
- [34] Johannes Buchner. *An Intuition for Physicists: Information Gain from Experiments*. Version 3. Aug. 26, 2022. DOI: [10.48550/arXiv.2205.00009](https://doi.org/10.48550/arXiv.2205.00009). arXiv: [2205.00009](https://arxiv.org/abs/2205.00009) [astro-ph, physics:cond-mat, physics:physics]. URL: <http://arxiv.org/abs/2205.00009> (visited on 05/29/2024). Pre-published.
- [35] Johannes Buchner. “Nested Sampling Methods”. In: *ArXiv210109675 Astro-Ph Stat* (July 13, 2021). arXiv: [2101.09675](https://arxiv.org/abs/2101.09675) [astro-ph, stat]. URL: <http://arxiv.org/abs/2101.09675> (visited on 10/19/2021).
- [36] Alessandra Buonanno et al. “Comparison of Post-Newtonian Templates for Compact Binary Inspiral Signals in Gravitational-Wave Detectors”. In: *Physical Review D* 80.8 (Oct. 29, 2009), p. 084043. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.80.084043](https://doi.org/10.1103/PhysRevD.80.084043). arXiv: [0907.0700](https://arxiv.org/abs/0907.0700). URL: <http://arxiv.org/abs/0907.0700> (visited on 03/15/2021).
- [37] Ollie Burke et al. *Mind the Gap: Addressing Data Gaps and Assessing Noise Mis-modeling in LISA*. Feb. 24, 2025. DOI: [10.48550/arXiv.2502.17426](https://doi.org/10.48550/arXiv.2502.17426). arXiv: [2502.17426](https://arxiv.org/abs/2502.17426) [gr-qc]. URL: <http://arxiv.org/abs/2502.17426> (visited on 02/26/2025). Pre-published.
- [38] Craig Cahillane et al. “Calibration Uncertainty for Advanced LIGO’s First and Second Observing Runs”. In: *Physical Review D* 96.10 (Nov. 8, 2017), p. 102001. DOI: [10.1103/PhysRevD.96.102001](https://doi.org/10.1103/PhysRevD.96.102001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.102001> (visited on 03/13/2026).

- [39] Ewan Cameron. “A Generalized Savage-Dickey Ratio”. In: *ArXiv13111292 Astro-Ph Stat* (Nov. 6, 2013). arXiv: [1311.1292 \[astro-ph, stat\]](https://arxiv.org/abs/1311.1292). URL: <http://arxiv.org/abs/1311.1292> (visited on 01/07/2021).
- [40] Priscilla Canizares et al. “Accelerated Gravitational Wave Parameter Estimation with Reduced Order Modeling”. In: *Physical Review Letters* 114 (Feb. 1, 2015), p. 071104. ISSN: 0031-9007. DOI: [10.1103/PhysRevLett.114.071104](https://doi.org/10.1103/PhysRevLett.114.071104). URL: <https://ui.adsabs.harvard.edu/abs/2015PhRvL.114g1104C> (visited on 06/06/2026).
- [41] Vitor Cardoso and Andrea Maselli. “Constraints on the Astrophysical Environment of Binaries with Gravitational-Wave Observations”. In: *Astron. Astrophys.* 644 (2020), A147. DOI: [10.1051/0004-6361/202037654](https://doi.org/10.1051/0004-6361/202037654).
- [42] Gregorio Carullo. “Black Hole Spectroscopy: Status Report”. In: *General Relativity and Gravitation* 57.5 (Apr. 29, 2025), p. 76. ISSN: 1572-9532. DOI: [10.1007/s10714-025-03408-y](https://doi.org/10.1007/s10714-025-03408-y). URL: <https://doi.org/10.1007/s10714-025-03408-y> (visited on 10/12/2025).
- [43] E. Chassande-Mottin and P. Flandrin. “On the Stationary Phase Approximation of Chirp Spectra”. In: *Proceedings of the IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis (Cat. No.98TH8380)*. International Symposium on Time-Frequency and Time-Scale Analysis. Pittsburgh, PA, USA: IEEE, 1998, pp. 117–120. ISBN: 978-0-7803-5073-1. DOI: [10.1109/TFSA.1998.721375](https://doi.org/10.1109/TFSA.1998.721375). URL: <http://ieeexplore.ieee.org/document/721375/> (visited on 08/24/2021).
- [44] Katerina Chatziioannou et al. “Noise Spectral Estimation Methods and Their Impact on Gravitational Wave Measurement of Compact Binary Mergers”. In: *Physical Review D* 100.10 (Nov. 5, 2019), p. 104004. DOI: [10.1103/PhysRevD.100.104004](https://doi.org/10.1103/PhysRevD.100.104004). URL: <https://link.aps.org/doi/10.1103/PhysRevD.100.104004> (visited on 03/04/2026).
- [45] Sushant Sharma Chaudhary et al. “Low-Latency Gravitational Wave Alert Products and Their Performance in Anticipation of the Fourth LIGO-Virgo-KAGRA Observing Run”. In: *Proceedings of the National Academy of Sciences* 121.18 (Apr. 30, 2024), e2316474121. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.2316474121](https://doi.org/10.1073/pnas.2316474121). arXiv: [2308.04545 \[astro-ph\]](https://arxiv.org/abs/2308.04545). URL: <http://arxiv.org/abs/2308.04545> (visited on 05/15/2024).
- [46] Anson Chen and Nathan K Johnson-McDaniel. “A Fast Frequency-Domain Expression for the Time-Dependent Detector Response of Ground-Based Gravitational-Wave Detectors to Compact Binary Signals”. In: (2024).
- [47] Danilo Chiaramello et al. “A Parametrized Model for Gravitational Waves from Eccentric, Precessing Binary Black Holes: Theory-Agnostic Tests of General Relativity with pTEOBResumS”. In: (Nov. 2025).

- [48] Qi Chu, Linqing Wen, and David Blair. “Scientific Benefit of Enlarging Gravitational Wave Detector Networks”. In: *J. Phys. Conf. Ser.* 363 (2012). Ed. by Mark Hannam et al., p. 012023. DOI: [10.1088/1742-6596/363/1/012023](https://doi.org/10.1088/1742-6596/363/1/012023).
- [49] Teagan A. Clarke et al. *A Universal Framework to Identify Eccentric Binary Mergers: GW200105 Case Study*. May 1, 2026. DOI: [10.48550/arXiv.2605.18742](https://doi.org/10.48550/arXiv.2605.18742). URL: <https://ui.adsabs.harvard.edu/abs/2026arXiv260518742C> (visited on 06/12/2026). Pre-published.
- [50] LIGO Scientific Collaboration, Virgo, and KAGRA. “Data Release for GW240925 and GW250207: Astrophysical Calibration of Gravitational-wave Detectors”. In: (Feb. 10, 2026). DOI: [10.5281/zenodo.18600070](https://doi.org/10.5281/zenodo.18600070). URL: <https://zenodo.org/records/18600070> (visited on 06/19/2026).
- [51] The LIGO Scientific Collaboration et al. “Properties of the Binary Neutron Star Merger GW170817”. In: *Physical Review X* 9.1 (Jan. 2, 2019), p. 011001. ISSN: 2160-3308. DOI: [10.1103/PhysRevX.9.011001](https://doi.org/10.1103/PhysRevX.9.011001). arXiv: [1805.11579](https://arxiv.org/abs/1805.11579). URL: <http://arxiv.org/abs/1805.11579> (visited on 03/10/2021).
- [52] The LIGO Scientific Collaboration et al. “The Basic Physics of the Binary Black Hole Merger GW150914”. In: *Annalen der Physik* 529.1–2 (Jan. 2017), p. 1600209. ISSN: 0003-3804, 1521-3889. DOI: [10.1002/andp.201600209](https://doi.org/10.1002/andp.201600209). arXiv: [1608.01940](https://arxiv.org/abs/1608.01940). URL: <http://arxiv.org/abs/1608.01940> (visited on 10/04/2021).
- [53] Roberto Cotesta et al. “Analysis of Ringdown Overtones in GW150914”. In: *Phys. Rev. Lett.* 129.11 (2022), p. 111102. DOI: [10.1103/PhysRevLett.129.111102](https://doi.org/10.1103/PhysRevLett.129.111102).
- [54] Michael Coughlin and Jan Harms. “Constraining the Gravitational Wave Energy Density of the Universe in the Range 0.1 Hz to 1 Hz Using the Apollo Seismic Array”. In: *Physical Review D* 90.10 (Nov. 3, 2014), p. 102001. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.90.102001](https://doi.org/10.1103/PhysRevD.90.102001). arXiv: [1409.4680](https://arxiv.org/abs/1409.4680) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/1409.4680> (visited on 07/07/2023).
- [55] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory 2nd Edition*. 2nd edition. Hoboken, N.J: Wiley-Interscience, July 18, 2006. 784 pp. ISBN: 978-0-471-24195-9.
- [56] Andrea Cozzumbo et al. “Opportunities and Limits of Lunar Gravitational-Wave Detection”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 382.2271 (Mar. 25, 2024), p. 20230066. DOI: [10.1098/rsta.2023.0066](https://doi.org/10.1098/rsta.2023.0066). URL: <https://royalsocietypublishing.org/doi/10.1098/rsta.2023.0066> (visited on 12/13/2024).
- [57] M. S. Cunha et al. “Astroseismology and Interferometry”. In: *The Astronomy and Astrophysics Review* 14.3 (Nov. 1, 2007), pp. 217–360. ISSN: 1432-0754. DOI: [10.1007/s00159-007-0007-0](https://doi.org/10.1007/s00159-007-0007-0). URL: <https://doi.org/10.1007/s00159-007-0007-0> (visited on 06/15/2026).

- [58] Thibault Damour. “Introductory Lectures on the Effective One Body Formalism”. In: *International Journal of Modern Physics A* 23.08 (Mar. 30, 2008), pp. 1130–1148. ISSN: 0217-751X, 1793-656X. DOI: [10.1142/S0217751X08039992](https://doi.org/10.1142/S0217751X08039992). arXiv: [0802.4047](https://arxiv.org/abs/0802.4047). URL: <http://arxiv.org/abs/0802.4047> (visited on 07/02/2021).
- [59] Thibault Damour and Alessandro Nagar. “A New Effective-One-Body Description of Coalescing Nonprecessing Spinning Black-Hole Binaries”. In: *Physical Review D* 90.4 (Aug. 13, 2014), p. 044018. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.90.044018](https://doi.org/10.1103/PhysRevD.90.044018). arXiv: [1406.6913](https://arxiv.org/abs/1406.6913). URL: <http://arxiv.org/abs/1406.6913> (visited on 08/31/2021).
- [60] Thibault Damour and Alessandro Nagar. “Faithful Effective-One-Body Waveforms of Small-Mass-Ratio Coalescing Black Hole Binaries”. In: *Physical Review D* 76.6 (Sept. 24, 2007), p. 064028. DOI: [10.1103/PhysRevD.76.064028](https://doi.org/10.1103/PhysRevD.76.064028). URL: <https://link.aps.org/doi/10.1103/PhysRevD.76.064028> (visited on 03/10/2026).
- [61] Maximilian Dax et al. *Group Equivariant Neural Posterior Estimation*. May 30, 2023. DOI: [10.48550/arXiv.2111.13139](https://doi.org/10.48550/arXiv.2111.13139). arXiv: [2111.13139](https://arxiv.org/abs/2111.13139) [cs]. URL: <http://arxiv.org/abs/2111.13139> (visited on 02/13/2025). Pre-published.
- [62] Maximilian Dax et al. *Neural Importance Sampling for Rapid and Reliable Gravitational-Wave Inference*. Oct. 11, 2022. arXiv: [2210.05686](https://arxiv.org/abs/2210.05686) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2210.05686> (visited on 11/14/2022). Pre-published.
- [63] Maximilian Dax et al. “Neural Importance Sampling for Rapid and Reliable Gravitational-Wave Inference”. In: *Phys. Rev. Lett.* 130.17 (2023), p. 171403. DOI: [10.1103/PhysRevLett.130.171403](https://doi.org/10.1103/PhysRevLett.130.171403).
- [64] Maximilian Dax et al. “Real-Time Gravitational-Wave Science with Neural Posterior Estimation”. In: *Physical Review Letters* 127.24 (Dec. 8, 2021), p. 241103. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.127.241103](https://doi.org/10.1103/PhysRevLett.127.241103). arXiv: [2106.12594](https://arxiv.org/abs/2106.12594) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2106.12594> (visited on 11/14/2022).
- [65] Marina De Amicis et al. “Dynamical Quasinormal Mode Excitation”. In: *Physical Review D* 113.2 (Jan. 22, 2026), p. 024048. DOI: [10.1103/bgjv-lvyd](https://doi.org/10.1103/bgjv-lvyd). URL: <https://link.aps.org/doi/10.1103/bgjv-lvyd> (visited on 04/02/2026).
- [66] James M. Dickey. “The Weighted Likelihood Ratio, Linear Hypotheses on Normal Location Parameters”. In: *The Annals of Mathematical Statistics* 42.1 (Feb. 1971), pp. 204–223. ISSN: 0003-4851, 2168-8990. DOI: [10.1214/aoms/1177693507](https://doi.org/10.1214/aoms/1177693507). URL: <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-42/issue-1/The-Weighted-Likelihood-Ratio-Linear-Hypotheses-on-Normal-Location-Parameters/10.1214/aoms/1177693507.full> (visited on 03/03/2026).

- [67] Ulyana Dupletsa et al. “Validating Prior-informed Fisher-matrix Analyses against GWTC Data”. In: *Physical Review D* 111.2 (Jan. 13, 2025), p. 024036. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.111.024036](https://doi.org/10.1103/PhysRevD.111.024036). arXiv: [2404.16103](https://arxiv.org/abs/2404.16103) [gr-qc]. URL: <http://arxiv.org/abs/2404.16103> (visited on 07/30/2025).
- [68] Dupletsa et al. “Gwfish: A Simulation Software to Evaluate Parameter-Estimation Capabilities of Gravitational-Wave Detector Networks”. In: *Astronomy and Computing* 42 (Jan. 1, 2023), p. 100671. ISSN: 2213-1337. DOI: [10.1016/j.ascom.2022.100671](https://doi.org/10.1016/j.ascom.2022.100671). URL: <https://www.sciencedirect.com/science/article/pii/S2213133722000853> (visited on 12/16/2022).
- [69] Freeman J. Dyson. “Seismic Response of the Earth to a Gravitational Wave in the 1-Hz Band”. In: *The Astrophysical Journal* 156 (May 1, 1969), p. 529. ISSN: 0004-637X. DOI: [10.1086/149986](https://doi.org/10.1086/149986). URL: <https://ui.adsabs.harvard.edu/abs/1969ApJ...156..529D> (visited on 03/15/2026).
- [70] Matthew C. Edwards, Renate Meyer, and Nelson Christensen. *Bayesian Semiparametric Power Spectral Density Estimation with Applications in Gravitational Wave Data Analysis*. Aug. 18, 2015. DOI: [10.48550/arXiv.1506.00185](https://doi.org/10.48550/arXiv.1506.00185). arXiv: [1506.00185](https://arxiv.org/abs/1506.00185). URL: <http://arxiv.org/abs/1506.00185> (visited on 10/28/2024). Pre-published.
- [71] Mattia Emma and Gregory Ashton. *RNLE: Residual Neural Likelihood Estimation and Its Application to Gravitational-Wave Astronomy*. Jan. 20, 2026. DOI: [10.48550/arXiv.2601.13857](https://doi.org/10.48550/arXiv.2601.13857). arXiv: [2601.13857](https://arxiv.org/abs/2601.13857) [gr-qc]. URL: <http://arxiv.org/abs/2601.13857> (visited on 03/03/2026). Pre-published.
- [72] Reed Essick, Salvatore Vitale, and Matthew Evans. “Frequency-Dependent Responses in Third Generation Gravitational-Wave Detectors”. In: *Physical Review D* 96.8 (Oct. 3, 2017), p. 084004. DOI: [10.1103/PhysRevD.96.084004](https://doi.org/10.1103/PhysRevD.96.084004). URL: <https://link.aps.org/doi/10.1103/PhysRevD.96.084004> (visited on 04/17/2024).
- [73] Stephen Fairhurst. “Triangulation of Gravitational Wave Sources with a Network of Detectors”. In: *New Journal of Physics* 11.12 (Dec. 2009), p. 123006. ISSN: 1367-2630. DOI: [10.1088/1367-2630/11/12/123006](https://doi.org/10.1088/1367-2630/11/12/123006). URL: <https://doi.org/10.1088/1367-2630/11/12/123006> (visited on 12/03/2025).
- [74] Will M Farr, Benjamin F. Farr, and T. Littenberg. *MODELLING CALIBRATION ERRORS IN CBC WAVEFORMS*. LIGO-T1400682. 2014. URL: <https://dcc.ligo.org/LIGO-T1400682/public>.
- [75] Marc Favata. “The Gravitational-Wave Memory from Eccentric Binaries”. In: *Physical Review D* 84.12 (Dec. 6, 2011), p. 124013. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.84.124013](https://doi.org/10.1103/PhysRevD.84.124013). arXiv: [1108.3121](https://arxiv.org/abs/1108.3121). URL: <http://arxiv.org/abs/1108.3121> (visited on 08/22/2021).

- [76] Lee Samuel Finn. “Detection, Measurement and Gravitational Radiation”. In: *Physical Review D* 46.12 (Dec. 15, 1992), pp. 5236–5249. ISSN: 0556-2821. DOI: [10.1103/PhysRevD.46.5236](https://doi.org/10.1103/PhysRevD.46.5236). arXiv: [gr-qc/9209010](https://arxiv.org/abs/gr-qc/9209010). URL: <http://arxiv.org/abs/gr-qc/9209010> (visited on 05/01/2023).
- [77] Eanna E. Flanagan and Tanja Hinderer. “Constraining Neutron Star Tidal Love Numbers with Gravitational Wave Detectors”. In: *Physical Review D* 77.2 (Jan. 22, 2008), p. 021502. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.77.021502](https://doi.org/10.1103/PhysRevD.77.021502). arXiv: [0709.1915](https://arxiv.org/abs/0709.1915) [astro-ph]. URL: <http://arxiv.org/abs/0709.1915> (visited on 06/26/2026).
- [78] Giulia Fumagalli et al. “Residual Eccentricity as a Systematic Uncertainty on the Formation Channels of Binary Black Holes”. In: *Physical Review D* 110.6 (Sept. 2024), p. 063012. ISSN: 1550-7998. DOI: [10.1103/PhysRevD.110.063012](https://doi.org/10.1103/PhysRevD.110.063012). URL: <https://ui.adsabs.harvard.edu/abs/2024PhRvD.110f3012F/abstract> (visited on 06/14/2026).
- [79] Rossella Gamba and Sebastiano Bernuzzi. “Resonant Tides in Binary Neutron Star Mergers: Analytical-numerical Relativity Study”. In: *Physical Review D* 107.4 (Feb. 8, 2023), p. 044014. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.107.044014](https://doi.org/10.1103/PhysRevD.107.044014). URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.044014> (visited on 06/14/2026).
- [80] Rossella Gamba, Sebastiano Bernuzzi, and Alessandro Nagar. “Fast, Faithful, Frequency-Domain Effective-One-Body Waveforms for Compact Binary Coalescences”. In: *ArXiv201200027 Astro-Ph Physicsgr-Qc* (Nov. 30, 2020). arXiv: [2012.00027](https://arxiv.org/abs/2012.00027) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2012.00027> (visited on 02/27/2021).
- [81] Rossella Gamba et al. “Revisiting GW150914 with a Non-Planar, Eccentric Waveform Model”. In: *Classical and Quantum Gravity* 42.17 (Sept. 2025), p. 175014. ISSN: 0264-9381. DOI: [10.1088/1361-6382/adfe50](https://doi.org/10.1088/1361-6382/adfe50). URL: <https://doi.org/10.1088/1361-6382/adfe50> (visited on 10/06/2025).
- [82] Cecilio García-Quirós et al. “IMRPhenomXHM: A Multi-Mode Frequency-Domain Model for the Gravitational Wave Signal from Non-Precessing Black-Hole Binaries”. In: *Physical Review D* 102.6 (Sept. 2, 2020), p. 064002. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.102.064002](https://doi.org/10.1103/PhysRevD.102.064002). arXiv: [2001.10914](https://arxiv.org/abs/2001.10914). URL: <http://arxiv.org/abs/2001.10914> (visited on 10/06/2021).
- [83] Davide Gerosa et al. “Which Is Which? Identification of the Two Compact Objects in Gravitational-Wave Binaries”. In: *Phys. Rev. Lett.* 134.12 (2025), p. 121402. DOI: [10.1103/PhysRevLett.134.121402](https://doi.org/10.1103/PhysRevLett.134.121402).
- [84] J. J. Giganti et al. *Lunar Surface Gravimeter Experiment*. Jan. 1, 1973. URL: <https://ntrs.nasa.gov/citations/19740010329> (visited on 03/15/2026).

- [85] Kirthinian Gobe et al. *Seismic Investigation in the Lunar Environment for Cosmic Events*. Zürich: ETH, 2025.
- [86] Eric Gourgoulhon. “3+1 Formalism and Bases of Numerical Relativity”. In: *ArXivgr-Qc0703035* (Mar. 6, 2007). arXiv: [gr-qc/0703035](https://arxiv.org/abs/gr-qc/0703035). URL: <http://arxiv.org/abs/gr-qc/0703035> (visited on 04/13/2021).
- [87] K. Grover et al. “Comparison of Gravitational Wave Detector Network Sky Localization Approximations”. In: *Physical Review D* 89.4 (Feb. 19, 2014), p. 042004. DOI: [10.1103/PhysRevD.89.042004](https://doi.org/10.1103/PhysRevD.89.042004). URL: <https://link.aps.org/doi/10.1103/PhysRevD.89.042004> (visited on 03/05/2026).
- [88] Nir Guttman et al. *Licence to Bin: Accurate and Scalable Inference for Binary Neutron Stars in Next-Generation Gravitational-Wave Detectors*. June 12, 2026. DOI: [10.48550/arXiv.2606.14197](https://doi.org/10.48550/arXiv.2606.14197). arXiv: [2606.14197](https://arxiv.org/abs/2606.14197) [astro-ph.HE]. URL: <http://arxiv.org/abs/2606.14197> (visited on 06/19/2026). Pre-published.
- [89] Jan Harms et al. “Lunar Gravitational-wave Antenna”. In: *The Astrophysical Journal* 910.1 (Mar. 2021), p. 1. ISSN: 0004-637X. DOI: [10.3847/1538-4357/abe5a7](https://doi.org/10.3847/1538-4357/abe5a7). URL: <https://doi.org/10.3847/1538-4357/abe5a7> (visited on 09/16/2021).
- [90] S. W. Hawking. “Gravitational Radiation from Colliding Black Holes”. In: *Physical Review Letters* 26.21 (May 24, 1971), pp. 1344–1346. DOI: [10.1103/PhysRevLett.26.1344](https://doi.org/10.1103/PhysRevLett.26.1344). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.26.1344> (visited on 03/13/2026).
- [91] M. Heideman, D. Johnson, and C. Burrus. “Gauss and the History of the Fast Fourier Transform”. In: *IEEE ASSP Magazine* 1.4 (Oct. 1984), pp. 14–21. ISSN: 1558-1284. DOI: [10.1109/MASSP.1984.1162257](https://doi.org/10.1109/MASSP.1984.1162257). URL: <https://ieeexplore.ieee.org/document/1162257> (visited on 03/11/2026).
- [92] Edward Higson et al. “Dynamic Nested Sampling: An Improved Algorithm for Parameter Estimation and Evidence Calculation”. In: *Statistics and Computing* 29.5 (Sept. 1, 2019), pp. 891–913. ISSN: 1573-1375. DOI: [10.1007/s11222-018-9844-0](https://doi.org/10.1007/s11222-018-9844-0). URL: <https://doi.org/10.1007/s11222-018-9844-0> (visited on 11/25/2024).
- [93] S Hild et al. “Sensitivity Studies for Third-Generation Gravitational Wave Observatories”. In: *Classical and Quantum Gravity* 28.9 (Apr. 2011), p. 094013. ISSN: 0264-9381. DOI: [10.1088/0264-9381/28/9/094013](https://doi.org/10.1088/0264-9381/28/9/094013). URL: <https://doi.org/10.1088/0264-9381/28/9/094013> (visited on 03/18/2026).
- [94] David W. Hogg and Daniel Foreman-Mackey. “Data Analysis Recipes: Using Markov Chain Monte Carlo*”. In: *The Astrophysical Journal Supplement Series* 236.1 (May 2018), p. 11. ISSN: 0067-0049. DOI: [10.3847/1538-4365/aab76e](https://doi.org/10.3847/1538-4365/aab76e). URL: <https://doi.org/10.3847/1538-4365/aab76e> (visited on 03/03/2026).

- [95] Francesco Iacovelli et al. *Forecasting the Detection Capabilities of Third-Generation Gravitational-Wave Detectors Using $\texttt{GFAST}$* . Sept. 27, 2022. DOI: [10.48550/arXiv.2207.02771](https://doi.org/10.48550/arXiv.2207.02771). arXiv: [2207.02771](https://arxiv.org/abs/2207.02771) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2207.02771> (visited on 10/24/2022). Pre-published.
- [96] Francesco Iacovelli et al. “Gravitational-Wave Parameter Estimation to the Moon and Back: Massive Binaries and the Case of GW231123”. In: *Phys. Rev. D* 113 (June 23, 2026), p. 123062. DOI: [10.1103/d2rh-btf9](https://doi.org/10.1103/d2rh-btf9). arXiv: [2512.09978](https://arxiv.org/abs/2512.09978) [gr-qc]. URL: <https://journals.aps.org/prd/abstract/10.1103/d2rh-btf9> (visited on 01/27/2026).
- [97] Francesco Iacovelli et al. *Supplement to “Forecasting the Detection Capabilities of Third-Generation Gravitational-Wave Detectors Using $\texttt{GFAST}$ ”: $\texttt{how To}$ $\texttt{GFAST}$* . July 20, 2022. arXiv: [2207.06910](https://arxiv.org/abs/2207.06910) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2207.06910> (visited on 09/09/2022). Pre-published.
- [98] Maximiliano Isi and Will M. Farr. *Analyzing Black-Hole Ringdowns*. July 13, 2021. arXiv: [2107.05609](https://arxiv.org/abs/2107.05609) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2107.05609> (visited on 07/11/2023). Pre-published.
- [99] Maximiliano Isi, Will M. Farr, and Katerina Chatziioannou. “Comparing Bayes Factors and Hierarchical Inference for Testing General Relativity with Gravitational Waves”. In: *ArXiv2204.10742 Astro-Ph Physicsgr-Qc* (Apr. 22, 2022). arXiv: [2204.10742](https://arxiv.org/abs/2204.10742) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2204.10742> (visited on 04/25/2022).
- [100] Justin Janquart et al. “Analyses of Overlapping Gravitational Wave Signals Using Hierarchical Subtraction and Joint Parameter Estimation”. In: *Monthly Notices of the Royal Astronomical Society* 523.2 (Aug. 1, 2023), pp. 1699–1710. ISSN: 0035-8711. DOI: [10.1093/mnras/stad1542](https://doi.org/10.1093/mnras/stad1542). URL: <https://doi.org/10.1093/mnras/stad1542> (visited on 03/18/2026).
- [101] Xisco Jimenez-Forteza et al. “Impact of High-Order Tidal Terms on Binary Neutron-Star Waveforms”. In: *Physical Review D* 98.12 (Dec. 13, 2018), p. 124014. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.98.124014](https://doi.org/10.1103/PhysRevD.98.124014). arXiv: [1807.08016](https://arxiv.org/abs/1807.08016) [gr-qc]. URL: <http://arxiv.org/abs/1807.08016> (visited on 06/15/2026).
- [102] Joris van Heijningen et al. *The Payload of the Lunar Gravitational-wave Antenna*. Feb. 1, 2023. arXiv: [2301.13685](https://arxiv.org/abs/2301.13685) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2301.13685> (visited on 02/02/2023). Pre-published.
- [103] Hannu Karttunen et al. “Celestial Mechanics”. In: *Fundamental Astronomy*. Ed. by Hannu Karttunen et al. Berlin, Heidelberg: Springer, 2017, pp. 123–140. ISBN: 978-3-662-53045-0. DOI: [10.1007/978-3-662-53045-0_6](https://doi.org/10.1007/978-3-662-53045-0_6). URL: https://doi.org/10.1007/978-3-662-53045-0_6 (visited on 04/21/2020).

- [104] Michael L. Katz et al. “FastEMRIWaveforms: New Tools for Millihertz Gravitational-Wave Data Analysis”. In: *Physical Review D* 104.6 (Sept. 17, 2021), p. 064047. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.104.064047](https://doi.org/10.1103/PhysRevD.104.064047). arXiv: [2104.04582](https://arxiv.org/abs/2104.04582) [gr-qc]. URL: <http://arxiv.org/abs/2104.04582> (visited on 01/23/2023).
- [105] D. E. Knuth. “Literate Programming”. In: *The Computer Journal* 27.2 (Jan. 1, 1984), pp. 97–111. ISSN: 0010-4620. DOI: [10.1093/comjnl/27.2.97](https://doi.org/10.1093/comjnl/27.2.97). URL: <https://doi.org/10.1093/comjnl/27.2.97> (visited on 03/05/2026).
- [106] Donald E. Knuth. “Two Notes on Notation”. Apr. 30, 1992. arXiv: [math/9205211](https://arxiv.org/abs/math/9205211). URL: <http://arxiv.org/abs/math/9205211> (visited on 03/04/2020).
- [107] Annalena Kofler et al. *Flexible Gravitational-Wave Parameter Estimation with Transformers*. Dec. 2, 2025. DOI: [10.48550/arXiv.2512.02968](https://doi.org/10.48550/arXiv.2512.02968). arXiv: [2512.02968](https://arxiv.org/abs/2512.02968) [gr-qc]. URL: <http://arxiv.org/abs/2512.02968> (visited on 03/04/2026). Pre-published.
- [108] Andrey Nikolaevich Kolmogorov. “Sulla Determinazione Empirica Di Una Legge Di Distribuzione”. In: 4 (1933), pp. 83–91.
- [109] Dong Lai. “Resonant Oscillations and Tidal Heating in Coalescing Binary Neutron Stars”. In: *Monthly Notices of the Royal Astronomical Society* 270.3 (Oct. 1994), pp. 611–629. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/270.3.611](https://doi.org/10.1093/mnras/270.3.611). arXiv: [astro-ph/9404062](https://arxiv.org/abs/astro-ph/9404062). URL: <http://arxiv.org/abs/astro-ph/9404062> (visited on 06/15/2026).
- [110] Jacob Lange, Richard O’Shaughnessy, and Monica Rizzo. “Rapid and Accurate Parameter Inference for Coalescing, Precessing Compact Binaries”. In: (May 2018).
- [111] Adam Lanman. *Aelanman/Lunarsky*. Dec. 24, 2024. URL: <https://github.com/aelanman/lunarsky> (visited on 03/11/2025).
- [112] Bo Liang and He Wang. “Recent Advances in Simulation-based Inference for Gravitational Wave Data Analysis”. In: *Astronomical Techniques and Instruments* 2.0 (2025), pp. 1–18. ISSN: 2097-3675, 2097-5880. DOI: [10.61977/ati2025020](https://doi.org/10.61977/ati2025020). arXiv: [2507.11192](https://arxiv.org/abs/2507.11192) [gr-qc]. URL: <http://arxiv.org/abs/2507.11192> (visited on 07/17/2025).
- [113] Virgo Collaboration and KAGRA Collaboration LIGO Scientific Collaboration. *GW250114 Discovery Paper Figure Scripts*. Sept. 10, 2025. DOI: [10.5281/zenodo.16877102](https://doi.org/10.5281/zenodo.16877102). URL: <https://zenodo.org/records/16877102> (visited on 06/21/2026).
- [114] LIGO Scientific Collaboration and Virgo Collaboration et al. “Observation of Gravitational Waves from a Binary Black Hole Merger”. In: *Physical Review Letters* 116.6 (Feb. 11, 2016), p. 061102. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.116.061102> (visited on 03/23/2020).
- [115] LIGO Scientific, Virgo, and KAGRA Collaborations. “GW240925 and GW250207: Astrophysical Calibration of Gravitational Wave Detectors”. In: *Physical Review Letters* (Apr. 30, 2026). ISSN: 0031-9007, 1079-7114. DOI: [10.1103/gzrj-mwv3](https://doi.org/10.1103/gzrj-mwv3). URL: <https://link.aps.org/doi/10.1103/gzrj-mwv3> (visited on 06/05/2026).

- [116] Frank Löffler et al. “The Einstein Toolkit: A Community Computational Infrastructure for Relativistic Astrophysics”. In: *Classical and Quantum Gravity* 29.11 (June 7, 2012), p. 115001. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/29/11/115001](https://doi.org/10.1088/0264-9381/29/11/115001). arXiv: [1111.3344](https://arxiv.org/abs/1111.3344). URL: <http://arxiv.org/abs/1111.3344> (visited on 08/27/2021).
- [117] *Lunar/LROC :: QuickMap*. URL: <https://tinyurl.com/quickmap-lgwa-position> (visited on 03/14/2026).
- [118] Ryan Lynch et al. “Observational Implications of Lowering the LIGO-Virgo Alert Threshold”. In: *The Astrophysical Journal* 861.2 (July 13, 2018), p. L24. ISSN: 2041-8213. DOI: [10.3847/2041-8213/aac9f9f](https://doi.org/10.3847/2041-8213/aac9f9f). arXiv: [1803.02880](https://arxiv.org/abs/1803.02880) [astro-ph]. URL: <http://arxiv.org/abs/1803.02880> (visited on 02/21/2023).
- [119] Michele Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. 1 edition. Oxford: Oxford University Press, Nov. 24, 2007. 576 pp. ISBN: 978-0-19-857074-5.
- [120] Sylvain Marsat and John G. Baker. *Fourier-Domain Modulations and Delays of Gravitational-Wave Signals*. June 27, 2018. DOI: [10.48550/arXiv.1806.10734](https://doi.org/10.48550/arXiv.1806.10734). arXiv: [1806.10734](https://arxiv.org/abs/1806.10734) [gr-qc]. URL: <http://arxiv.org/abs/1806.10734> (visited on 12/07/2022). Pre-published.
- [121] Sylvain Marsat, John G. Baker, and Tito Dal Canton. “Exploring the Bayesian Parameter Estimation of Binary Black Holes with LISA”. In: *Physical Review D* 103.8 (Apr. 15, 2021), p. 083011. DOI: [10.1103/PhysRevD.103.083011](https://doi.org/10.1103/PhysRevD.103.083011). URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.083011> (visited on 01/15/2025).
- [122] H. J. Melosh. “New Approaches to the Moon’s Isotopic Crisis”. In: *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences* 372.2024 (Sept. 13, 2014), p. 20130168. ISSN: 1364-503X. DOI: [10.1098/rsta.2013.0168](https://doi.org/10.1098/rsta.2013.0168). PMID: [25114301](https://pubmed.ncbi.nlm.nih.gov/25114301/).
- [123] Francesco Messina et al. “Quasi-5.5PN TaylorF2 Approximant for Compact Binaries: Point-Mass Phasing and Impact on the Tidal Polarizability Inference”. In: *Physical Review D* 99.12 (June 28, 2019), p. 124051. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.99.124051](https://doi.org/10.1103/PhysRevD.99.124051). arXiv: [1904.09558](https://arxiv.org/abs/1904.09558). URL: <http://arxiv.org/abs/1904.09558> (visited on 03/17/2021).
- [124] C.W. Misner, K.S. Thorne, and J.A. Wheeler. *Gravitation*. W.H. Freeman & Co., 1973. ISBN: 978-0-691-17779-3.
- [125] Christopher J. Moore, Robert H. Cole, and Christopher P. L. Berry. “Gravitational-Wave Sensitivity Curves”. In: *Classical and Quantum Gravity* 32.1 (Jan. 8, 2015), p. 015014. ISSN: 0264-9381, 1361-6382. DOI: [10.1088/0264-9381/32/1/015014](https://doi.org/10.1088/0264-9381/32/1/015014). arXiv: [1408.0740](https://arxiv.org/abs/1408.0740). URL: <http://arxiv.org/abs/1408.0740> (visited on 04/19/2020).

- [126] Soichiro Morisaki et al. *Rapid Localization and Inference on Compact Binary Coalescences with the Advanced LIGO-Virgo-KAGRA Gravitational-Wave Detector Network*. July 25, 2023. arXiv: [2307.13380](#) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2307.13380> (visited on 07/26/2023). Pre-published.
- [127] Alessandro Nagar and Piero Retteno. “Rush the Inspiral: Efficient Effective One Body Time-Domain Gravitational Waveforms”. In: *ArXiv180503891 Gr-Qc* (Nov. 25, 2019). DOI: [10.1103/PhysRevD.99.021501](#). arXiv: [1805.03891](#) [gr-qc]. URL: <http://arxiv.org/abs/1805.03891> (visited on 02/27/2021).
- [128] Alessandro Nagar et al. “Effective-One-Body Waveform Model for Noncircularized, Planar, Coalescing Black Hole Binaries: The Importance of Radiation Reaction”. In: *Physical Review D* 110.8 (Oct. 1, 2024), p. 084001. DOI: [10.1103/PhysRevD.110.84001](#). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.84001> (visited on 03/10/2026).
- [129] Alessandro Nagar et al. “Time-Domain Effective-One-Body Gravitational Waveforms for Coalescing Compact Binaries with Nonprecessing Spins, Tides and Self-Spin Effects”. In: *ArXiv180601772 Gr-Qc* (Oct. 28, 2018). DOI: [10.1103/PhysRevD.98.104052](#). arXiv: [1806.01772](#) [gr-qc]. URL: <http://arxiv.org/abs/1806.01772> (visited on 02/27/2021).
- [130] Harsh Narola et al. *Relative Binning for Complete Gravitational-Wave Parameter Estimation with Higher-Order Modes and Precession, and Applications to Lensing and Third-Generation Detectors*. Aug. 23, 2023. arXiv: [2308.12140](#) [gr-qc]. URL: <http://arxiv.org/abs/2308.12140> (visited on 08/28/2023). Pre-published.
- [131] Atsushi Nishizawa et al. “Probing Non-Tensorial Polarizations of Stochastic Gravitational-Wave Backgrounds with Ground-Based Laser Interferometers”. In: *Physical Review D* 79.8 (Apr. 16, 2009), p. 082002. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.79.082002](#). arXiv: [0903.0528](#) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/0903.0528> (visited on 02/08/2023).
- [132] Alexander H. Nitz et al. “Rapid Detection of Gravitational Waves from Compact Binary Mergers with PyCBC Live”. In: *Physical Review D* 98 (July 1, 2018), p. 024050. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.98.024050](#). URL: <https://ui.adsabs.harvard.edu/abs/2018PhRvD..98b4050N> (visited on 02/20/2023).
- [133] Serguei Ossokine et al. “Multipolar Effective-One-Body Waveforms for Precessing Binary Black Holes: Construction and Validation”. In: *Physical Review D* 102 (Aug. 1, 2020), p. 044055. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.102.044055](#). URL: <https://ui.adsabs.harvard.edu/abs/2020PhRvD.102d4055O> (visited on 03/09/2026).

- [134] H. Padé. “Sur la représentation approchée d’une fonction par des fractions rationnelles”. In: *Annales scientifiques de l’École Normale Supérieure* 9 (1892), pp. 3–93. ISSN: 1873-2151. DOI: [10.24033/asens.378](https://doi.org/10.24033/asens.378). URL: https://www.numdam.org/item/ASENS_1892_3_9__S3_0/ (visited on 03/10/2026).
- [135] George Papamakarios and Iain Murray. “Fast ϵ -Free Inference of Simulation Models with Bayesian Conditional Density Estimation”. In: *Advances in Neural Information Processing Systems*. Vol. 29. Curran Associates, Inc., 2016. URL: <https://proceedings.neurips.cc/paper/2016/hash/6aca97005c68f1206823815f66102863-Abstract.html> (visited on 03/03/2026).
- [136] George Papamakarios et al. *Normalizing Flows for Probabilistic Modeling and Inference*. Apr. 8, 2021. DOI: [10.48550/arXiv.1912.02762](https://doi.org/10.48550/arXiv.1912.02762). arXiv: [1912.02762](https://arxiv.org/abs/1912.02762) [stat]. URL: <http://arxiv.org/abs/1912.02762> (visited on 03/03/2026). Pre-published.
- [137] Vilfredo Pareto. *Manuale di economia politica con una introduzione alla scienza sociale*. Milano : Societa Editrice Libreria, 1919. 600 pp. URL: <http://archive.org/details/manualedieconomi00pareuoft> (visited on 06/11/2026).
- [138] P. C. Peters and J. Mathews. “Gravitational Radiation from Point Masses in a Keplerian Orbit”. In: *Physical Review* 131.1 (July 1, 1963), pp. 435–440. ISSN: 0031-899X. DOI: [10.1103/PhysRev.131.435](https://doi.org/10.1103/PhysRev.131.435). URL: <https://link.aps.org/doi/10.1103/PhysRev.131.435> (visited on 04/04/2021).
- [139] Aleksandr Petrosyan and William James Handley. *SuperNest: Accelerated Nested Sampling Applied to Astrophysics and Cosmology*. Dec. 4, 2022. arXiv: [2212.01760](https://arxiv.org/abs/2212.01760) [astro-ph, physics:physics]. URL: <http://arxiv.org/abs/2212.01760> (visited on 12/06/2022). Pre-published.
- [140] Hugo Pfister et al. “Observable Gravitational Waves from Tidal Disruption Events and Their Electromagnetic Counterpart”. In: *Monthly Notices of the Royal Astronomical Society* 510.2 (Feb. 21, 2022), pp. 2025–2040. ISSN: 0035-8711. DOI: [10.1093/mnras/stab3387](https://doi.org/10.1093/mnras/stab3387). URL: <https://doi.org/10.1093/mnras/stab3387> (visited on 06/23/2026).
- [141] Khun Sang Phukon, Patricia Schmidt, and Geraint Pratten. “Geometric Template Bank for the Detection of Spinning Low-Mass Compact Binaries with Moderate Orbital Eccentricity”. In: *Phys. Rev. D* 111.4 (2025), p. 043040. DOI: [10.1103/PhysRevD.111.043040](https://doi.org/10.1103/PhysRevD.111.043040).
- [142] F. A. E. Pirani. “On the Physical Significance of the Riemann Tensor”. In: *Acta Physica Polonica* 15 (Jan. 1, 1956), pp. 389–405. URL: <https://ui.adsabs.harvard.edu/abs/1956AcPP...15..389P> (visited on 03/15/2026).

- [143] Adam Pound et al. “Second-Order Self-Force Calculation of Gravitational Binding Energy in Compact Binaries”. In: *Physical Review Letters* 124.2 (Jan. 14, 2020), p. 021101. DOI: [10.1103/PhysRevLett.124.021101](https://doi.org/10.1103/PhysRevLett.124.021101). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.021101> (visited on 06/24/2026).
- [144] Geraint Pratten et al. “Computationally Efficient Models for the Dominant and Subdominant Harmonic Modes of Precessing Binary Black Holes”. In: *Physical Review D* 103.10 (May 25, 2021), p. 104056. DOI: [10.1103/PhysRevD.103.104056](https://doi.org/10.1103/PhysRevD.103.104056). URL: <https://link.aps.org/doi/10.1103/PhysRevD.103.104056> (visited on 03/08/2026).
- [145] Frans Pretorius. “Evolution of Binary Black Hole Spacetimes”. In: *Physical Review Letters* 95.12 (Sept. 14, 2005), p. 121101. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.95.121101](https://doi.org/10.1103/PhysRevLett.95.121101). arXiv: [gr-qc/0507014](https://arxiv.org/abs/gr-qc/0507014). URL: <http://arxiv.org/abs/gr-qc/0507014> (visited on 04/13/2021).
- [146] M Punturo et al. “The Einstein Telescope: A Third-Generation Gravitational Wave Observatory”. In: *Classical and Quantum Gravity* 27.19 (Sept. 2010), p. 194002. ISSN: 0264-9381. DOI: [10.1088/0264-9381/27/19/194002](https://doi.org/10.1088/0264-9381/27/19/194002). URL: <https://doi.org/10.1088/0264-9381/27/19/194002> (visited on 03/18/2026).
- [147] Tania Regimbau et al. “Mock Data Challenge for the Einstein Gravitational-Wave Telescope”. In: *Physical Review D* 86.12 (Dec. 3, 2012), p. 122001. DOI: [10.1103/PhysRevD.86.122001](https://doi.org/10.1103/PhysRevD.86.122001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.122001> (visited on 03/04/2026).
- [148] David Reitze et al. “Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO”. In: *Bull. Am. Astron. Soc.* 51.7 (2019), p. 035.
- [149] Justin L. Ripley et al. “A Constraint on the Dissipative Tidal Deformability of Neutron Stars”. In: *Nature Astron.* 8.10 (2024), pp. 1277–1283. DOI: [10.1038/s41550-024-02323-7](https://doi.org/10.1038/s41550-024-02323-7).
- [150] Carl L. Rodriguez et al. “Inadequacies of the Fisher Information Matrix in Gravitational-Wave Parameter Estimation”. In: *Physical Review D* 88.8 (Oct. 9, 2013), p. 084013. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.88.084013](https://doi.org/10.1103/PhysRevD.88.084013). arXiv: [1308.1397](https://arxiv.org/abs/1308.1397) [astro-ph]. URL: <http://arxiv.org/abs/1308.1397> (visited on 01/25/2024).
- [151] I. M. Romero-Shaw et al. “Bayesian Inference for Compact Binary Coalescences with BILBY: Validation and Application to the First LIGO–Virgo Gravitational-Wave Transient Catalogue”. In: *Monthly Notices of the Royal Astronomical Society* 499.3 (Oct. 29, 2020), pp. 3295–3319. ISSN: 0035-8711, 1365-2966. DOI: [10.1093/mnras/staa2850](https://doi.org/10.1093/mnras/staa2850). arXiv: [2006.00714](https://arxiv.org/abs/2006.00714) [astro-ph]. URL: <http://arxiv.org/abs/2006.00714> (visited on 05/30/2025).

- [152] Isobel M. Romero-Shaw, Eric Thrane, and Paul D. Lasky. “When Models Fail: An Introduction to Posterior Predictive Checks and Model Misspecification in Gravitational-Wave Astronomy”. In: *Publications of the Astronomical Society of Australia* 39 (2022), e025. ISSN: 1323-3580, 1448-6083. DOI: [10.1017/pasa.2022.24](https://doi.org/10.1017/pasa.2022.24). arXiv: [2202.05479](https://arxiv.org/abs/2202.05479) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2202.05479> (visited on 02/24/2024).
- [153] S. Ronchini et al. “Perspectives for Multimessenger Astronomy with the next Generation of Gravitational-Wave Detectors and High-Energy Satellites”. In: *Astronomy & Astrophysics* 665 (Sept. 1, 2022), A97. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202243705](https://doi.org/10.1051/0004-6361/202243705). URL: <https://www.aanda.org/articles/aa/abs/2022/09/aa43705-22/aa43705-22.html> (visited on 03/18/2026).
- [154] Javier Roulet and Tejaswi Venumadhav. *Inferring Binary Properties from Gravitational Wave Signals*. Feb. 17, 2024. DOI: [10.1146/annurev-nucl-121423-100725](https://doi.org/10.1146/annurev-nucl-121423-100725). arXiv: [2402.11439](https://arxiv.org/abs/2402.11439) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/2402.11439> (visited on 02/20/2024). Pre-published.
- [155] F. Santoliquido et al. “Comparing Next-Generation Detector Configurations for High-Redshift Gravitational Wave Sources with Neural Posterior Estimation”. In: *Astronomy & Astrophysics* 708 (Apr. 1, 2026), A175. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202558828](https://doi.org/10.1051/0004-6361/202558828). URL: <https://www.aanda.org/articles/aa/abs/2026/04/aa58828-25/aa58828-25.html> (visited on 05/20/2026).
- [156] Filippo Santoliquido et al. “Binary Black Hole Mergers from Population III Stars: Uncertainties from Star Formation and Binary Star Properties”. In: *Monthly Notices of the Royal Astronomical Society* 524.1 (Sept. 1, 2023), pp. 307–324. ISSN: 0035-8711. DOI: [10.1093/mnras/stad1860](https://doi.org/10.1093/mnras/stad1860). URL: <https://doi.org/10.1093/mnras/stad1860> (visited on 03/18/2026).
- [157] Filippo Santoliquido et al. “Classifying Binary Black Holes from Population III Stars with the Einstein Telescope: A Machine-Learning Approach”. In: *Astronomy & Astrophysics* 690 (Oct. 2024), A362. ISSN: 0004-6361, 1432-0746. DOI: [10.1051/0004-6361/202450381](https://doi.org/10.1051/0004-6361/202450381). arXiv: [2404.10048](https://arxiv.org/abs/2404.10048) [astro-ph]. URL: <http://arxiv.org/abs/2404.10048> (visited on 07/30/2025).
- [158] Filippo Santoliquido et al. *Comparing Next-Generation Detector Configurations for High-Redshift Gravitational Wave Sources with Neural Posterior Estimation*. Dec. 23, 2025. DOI: [10.48550/arXiv.2512.20699](https://doi.org/10.48550/arXiv.2512.20699). arXiv: [2512.20699](https://arxiv.org/abs/2512.20699) [gr-qc]. URL: <http://arxiv.org/abs/2512.20699> (visited on 03/02/2026). Pre-published.
- [159] Filippo Santoliquido et al. “Fast and Accurate Parameter Estimation of High-Redshift Sources with the Einstein Telescope”. In: *Phys. Rev. D* 112.10 (2025), p. 103015. DOI: [10.1103/wflk-p5cl](https://doi.org/10.1103/wflk-p5cl).

- [160] Matteo Serra. *LUSATIA IS OFFICIALLY THE THIRD CANDIDATE SITE FOR THE EINSTEIN TELESCOPE*. Einstein Telescope Italia. Oct. 21, 2025. URL: <http://www.einstein-telescope.it/en/2025/10/21/lusatia-is-officially-the-third-candidate-site-for-the-einstein-telescope/> (visited on 03/18/2026).
- [161] Md Arif Shaikh et al. “Defining Eccentricity for Gravitational Wave Astronomy”. In: *Phys. Rev. D* 108.10 (2023), p. 104007. DOI: [10.1103/PhysRevD.108.104007](https://doi.org/10.1103/PhysRevD.108.104007).
- [162] Neha Sharma, Aditya Vijaykumar, and Prayush Kumar. *Rapid Inference of Gravitational-Wave Signals in the Time Domain Using a Heterodyned Likelihood*. Jan. 16, 2026. DOI: [10.48550/arXiv.2601.11239](https://doi.org/10.48550/arXiv.2601.11239). arXiv: [2601.11239 \[gr-qc\]](https://arxiv.org/abs/2601.11239). URL: <http://arxiv.org/abs/2601.11239> (visited on 01/21/2026). Pre-published.
- [163] Leo P. Singer and Larry R. Price. “Rapid Bayesian Position Reconstruction for Gravitational-Wave Transients”. In: *Physical Review D* 93.2 (Jan. 14, 2016), p. 024013. DOI: [10.1103/PhysRevD.93.024013](https://doi.org/10.1103/PhysRevD.93.024013). URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.024013> (visited on 11/21/2022).
- [164] Devinderjit Sivia and John Skilling. *Data Analysis: A Bayesian Tutorial*. Oxford University Press, June 2006. 259 pp. ISBN: 978-0-19-856831-5. Google Books: [IYMSDAAAQBAJ](https://books.google.com/books?id=IYMSDAAAQBAJ).
- [165] John Skilling. “Nested Sampling”. In: *AIP Conference Proceedings*. BAYESIAN INFERENCE AND MAXIMUM ENTROPY METHODS IN SCIENCE AND ENGINEERING: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering. Vol. 735. Garching (Germany): AIP, 2004, pp. 395–405. DOI: [10.1063/1.1835238](https://doi.org/10.1063/1.1835238). URL: <https://pubs.aip.org/aip/acp/article/735/1/395-405/748716> (visited on 11/25/2024).
- [166] John Skilling. “Nested Sampling for General Bayesian Computation”. In: *Bayesian Analysis* 1.4 (Dec. 2006), pp. 833–859. ISSN: 1936-0975, 1931-6690. DOI: [10.1214/06-BA127](https://doi.org/10.1214/06-BA127). URL: <https://projecteuclid.org/journals/bayesian-analysis/volume-1/issue-4/Nested-sampling-for-general-Bayesian-computation/10.1214/06-BA127.full> (visited on 09/15/2021).
- [167] N. Smirnov. “Table for Estimating the Goodness of Fit of Empirical Distributions”. In: *The Annals of Mathematical Statistics* 19.2 (June 1948), pp. 279–281. ISSN: 0003-4851, 2168-8990. DOI: [10.1214/aoms/1177730256](https://doi.org/10.1214/aoms/1177730256). URL: <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-19/issue-2/Table-for-Estimating-the-Goodness-of-Fit-of-Empirical-Distributions/10.1214/aoms/1177730256.full> (visited on 03/08/2026).
- [168] Joshua S. Speagle. “DYNESTY: A Dynamic Nested Sampling Package for Estimating Bayesian Posteriors and Evidences”. In: *Monthly Notices of the Royal Astronomical Society* 493 (Apr. 1, 2020), pp. 3132–3158. ISSN: 0035-8711. DOI: [10.1093/mnras/staa278](https://doi.org/10.1093/mnras/staa278). URL: <https://ui.adsabs.harvard.edu/abs/2020MNRAS.493.3132S> (visited on 11/22/2024).

- [169] Lorenzo Speri et al. *Quantifying the Scientific Potential of Intermediate and Extreme Mass Ratio Inspirals with the Laser Interferometer Space Antenna*. Version 1. Mar. 17, 2026. DOI: [10.48550/arXiv.2603.17072](https://doi.org/10.48550/arXiv.2603.17072). arXiv: [2603.17072](https://arxiv.org/abs/2603.17072) [astro-ph.IM]. URL: <http://arxiv.org/abs/2603.17072> (visited on 06/24/2026). Pre-published.
- [170] Seth Stein and Michael Wysession. *An Introduction to Seismology, Earthquakes, and Earth Structure*. John Wiley & Sons, Apr. 1, 2009. 515 pp. ISBN: 978-1-4443-1131-0.
- [171] Jan Steinhoff et al. “Spin Effects on Neutron Star Fundamental-Mode Dynamical Tides: Phenomenology and Comparison to Numerical Simulations”. In: *Physical Review Research* 3.3 (Aug. 9, 2021), p. 033129. ISSN: 2643-1564. DOI: [10.1103/PhysRevResearch.3.033129](https://doi.org/10.1103/PhysRevResearch.3.033129). arXiv: [2103.06100](https://arxiv.org/abs/2103.06100) [gr-qc]. URL: <http://arxiv.org/abs/2103.06100> (visited on 06/15/2026).
- [172] János Takátsy, Bence Kocsis, and Péter Kovács. “Observability of Dynamical Tides in Merging Eccentric Neutron Star Binaries”. In: *Physical Review D* 110.10 (Nov. 26, 2024), p. 103043. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.110.103043](https://doi.org/10.1103/PhysRevD.110.103043). arXiv: [2407.17560](https://arxiv.org/abs/2407.17560) [astro-ph.HE]. URL: <http://arxiv.org/abs/2407.17560> (visited on 06/15/2026).
- [173] Colm Talbot et al. “Inference with Finite Time Series: II. The Window Strikes Back”. In: *Classical and Quantum Gravity* 42.23 (Dec. 2025), p. 235023. ISSN: 0264-9381. DOI: [10.1088/1361-6382/ae1ac7](https://doi.org/10.1088/1361-6382/ae1ac7). URL: <https://doi.org/10.1088/1361-6382/ae1ac7> (visited on 04/20/2026).
- [174] Colm Talbot et al. “Inference with Finite Time Series: Observing the Gravitational Universe through Windows”. In: *Physical Review Research* 3.4 (Oct. 18, 2021), p. 043049. DOI: [10.1103/PhysRevResearch.3.043049](https://doi.org/10.1103/PhysRevResearch.3.043049). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.3.043049> (visited on 12/16/2022).
- [175] Nicola Tamanini et al. “Peculiar Acceleration of Stellar-Origin Black Hole Binaries: Measurement and Biases with LISA”. In: *Physical Review D* 101.6 (Mar. 3, 2020), p. 063002. DOI: [10.1103/PhysRevD.101.063002](https://doi.org/10.1103/PhysRevD.101.063002). URL: <https://link.aps.org/doi/10.1103/PhysRevD.101.063002> (visited on 03/17/2026).
- [176] Rodrigo Tenorio and Davide Gerosa. *SFTs: A Scalable Data-Analysis Framework for Long-Duration Gravitational-Wave Signals*. Feb. 17, 2025. DOI: [10.48550/arXiv.2502.11823](https://doi.org/10.48550/arXiv.2502.11823). arXiv: [2502.11823](https://arxiv.org/abs/2502.11823) [gr-qc]. URL: <http://arxiv.org/abs/2502.11823> (visited on 05/12/2025). Pre-published.
- [177] The LIGO Scientific Collaboration et al. *GWTC-4.0: Updating the Gravitational-Wave Transient Catalog with Observations from the First Part of the Fourth LIGO-Virgo-KAGRA Observing Run*. Aug. 1, 2025. DOI: [10.48550/arXiv.2508.18082](https://doi.org/10.48550/arXiv.2508.18082). URL: <https://ui.adsabs.harvard.edu/abs/2025arXiv250818082T> (visited on 03/07/2026). Pre-published.

- [178] Eric Thrane and Colm Talbot. “An Introduction to Bayesian Inference in Gravitational-Wave Astronomy: Parameter Estimation, Model Selection, and Hierarchical Models”. In: *Publications of the Astronomical Society of Australia* 36 (Jan. 2019), e010. ISSN: 1323-3580, 1448-6083. DOI: [10.1017/pasa.2019.2](https://doi.org/10.1017/pasa.2019.2). URL: <https://www.cambridge.org/core/journals/publications-of-the-astronomical-society-of-australia/article/an-introduction-to-bayesian-inference-in-gravitational-wave-astronomy-parameter-estimation-model-selection-and-hierarchical-models/D459F61D8C37F0BEF86D60F42A418304> (visited on 02/07/2026).
- [179] Manuel Tiglio and Aarón Villanueva. “Reduced Order and Surrogate Models for Gravitational Waves”. In: *Living Reviews in Relativity* 25.1 (Apr. 26, 2022), p. 2. ISSN: 1433-8351. DOI: [10.1007/s41114-022-00035-w](https://doi.org/10.1007/s41114-022-00035-w). URL: <https://doi.org/10.1007/s41114-022-00035-w> (visited on 03/18/2026).
- [180] Jacopo Tissino. “Machine Learning for Gravitational Waveforms from Binary Neutron Star Mergers”. Padua, Italy, 2021. 100 pp.
- [181] Jacopo Tissino and Filippo Santoliquido. *The Geometry of Lunar Gravitational Wave Detection: Data Release*. Zenodo, June 2, 2026. DOI: [10.5281/zenodo.20510354](https://doi.org/10.5281/zenodo.20510354). URL: <https://zenodo.org/records/20510354> (visited on 06/18/2026).
- [182] Jacopo Tissino et al. “Combining Effective-One-Body Accuracy and Reduced-Order-Quadrature Speed for Binary Neutron Star Merger Parameter Estimation with Machine Learning”. In: *Phys. Rev. D* 107.8 (2023), p. 084037. DOI: <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.107.084037>. arXiv: [2210.15684](https://arxiv.org/abs/2210.15684) [gr-qc]. URL: <http://arxiv.org/abs/2210.15684> (visited on 11/04/2022).
- [183] Jacopo Tissino et al. *The Geometry of Lunar Gravitational Wave Detection*. June 3, 2026. DOI: [10.48550/arXiv.2606.04918](https://doi.org/10.48550/arXiv.2606.04918). arXiv: [2606.04918](https://arxiv.org/abs/2606.04918) [gr-qc]. URL: <http://arxiv.org/abs/2606.04918> (visited on 06/04/2026). Pre-published.
- [184] Martina Toscani et al. “Gravitational Waves from Tidal Disruption Events: An Open and Comprehensive Catalog”. In: *Monthly Notices of the Royal Astronomical Society* 510.1 (Feb. 11, 2022), pp. 992–1001. ISSN: 0035-8711. DOI: [10.1093/mnras/stab3384](https://doi.org/10.1093/mnras/stab3384). URL: <https://doi.org/10.1093/mnras/stab3384> (visited on 06/23/2026).
- [185] Alexandre Toubiana et al. “Parameter Estimation of Stellar-Mass Black Hole Binaries with LISA”. In: *Physical Review D* 102.12 (Dec. 16, 2020), p. 124037. DOI: [10.1103/PhysRevD.102.124037](https://link.aps.org/doi/10.1103/PhysRevD.102.124037). URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.124037> (visited on 12/03/2025).
- [186] Michele Vallisneri. “Use and Abuse of the Fisher Information Matrix in the Assessment of Gravitational-Wave Parameter-Estimation Prospects”. In: *Physical Review D* 77.4 (Feb. 4, 2008), p. 042001. DOI: [10.1103/PhysRevD.77.042001](https://link.aps.org/doi/10.1103/PhysRevD.77.042001). URL: <https://link.aps.org/doi/10.1103/PhysRevD.77.042001> (visited on 05/01/2023).

- [187] Vijay Varma et al. “Surrogate Models for Precessing Binary Black Hole Simulations with Unequal Masses”. In: *Physical Review Research* 1.3 (Oct. 10, 2019), p. 033015. DOI: [10.1103/PhysRevResearch.1.033015](https://doi.org/10.1103/PhysRevResearch.1.033015). URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.1.033015> (visited on 03/05/2026).
- [188] David Vartanyan. “Gravitational-Wave Signature of Core-Collapse Supernovae”. In: *Physical Review D* 107.10 (2023). DOI: [10.1103/PhysRevD.107.103015](https://doi.org/10.1103/PhysRevD.107.103015).
- [189] John Veitch and W. Del Pozzo. *Analytic Marginalisation of Phase Parameter*. 2013. URL: <https://dcc.ligo.org/LIGO-T1300326/public> (visited on 02/07/2026).
- [190] Leslie Wade et al. “Systematic and Statistical Errors in a Bayesian Approach to the Estimation of the Neutron-Star Equation of State Using Advanced Gravitational Wave Detectors”. In: *Phys. Rev. D* 89.10 (2014), p. 103012. DOI: [10.1103/PhysRevD.89.103012](https://doi.org/10.1103/PhysRevD.89.103012).
- [191] J. Weber. “Detection and Generation of Gravitational Waves”. In: *Physical Review* 117.1 (Jan. 1, 1960), pp. 306–313. DOI: [10.1103/PhysRev.117.306](https://doi.org/10.1103/PhysRev.117.306). URL: <https://link.aps.org/doi/10.1103/PhysRev.117.306> (visited on 02/16/2023).
- [192] Linqing Wen and Yanbei Chen. “Geometrical Expression for the Angular Resolution of a Network of Gravitational-Wave Detectors”. In: *Physical Review D* 81.8 (Apr. 8, 2010), p. 082001. ISSN: 1550-7998, 1550-2368. DOI: [10.1103/PhysRevD.81.082001](https://doi.org/10.1103/PhysRevD.81.082001). arXiv: [1003.2504](https://arxiv.org/abs/1003.2504) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/1003.2504> (visited on 03/24/2023).
- [193] P. Whittle. “Curve and Periodogram Smoothing”. In: *Journal of the Royal Statistical Society Series B: Statistical Methodology* 19.1 (Jan. 1, 1957), pp. 38–47. ISSN: 1369-7412, 1467-9868. DOI: [10.1111/j.2517-6161.1957.tb00242.x](https://doi.org/10.1111/j.2517-6161.1957.tb00242.x). URL: <https://academic.oup.com/jrsssb/article/19/1/38/7026758> (visited on 02/11/2025).
- [194] Jonas Wildberger et al. “Adapting to Noise Distribution Shifts in Flow-Based Gravitational-Wave Inference”. In: *Physical Review D* 107.8 (Apr. 26, 2023), p. 084046. DOI: [10.1103/PhysRevD.107.084046](https://doi.org/10.1103/PhysRevD.107.084046). URL: <https://link.aps.org/doi/10.1103/PhysRevD.107.084046> (visited on 03/04/2026).
- [195] Michael J. Williams, John Veitch, and Chris Messenger. “Nested Sampling with Normalising Flows for Gravitational-Wave Inference”. In: *Physical Review D* 103.10 (May 5, 2021), p. 103006. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.103.103006](https://doi.org/10.1103/PhysRevD.103.103006). arXiv: [2102.11056](https://arxiv.org/abs/2102.11056) [gr-qc]. URL: <http://arxiv.org/abs/2102.11056> (visited on 10/20/2025).
- [196] Barak Zackay, Liang Dai, and Tejaswi Venumadhav. “Relative Binning and Fast Likelihood Evaluation for Gravitational Wave Parameter Estimation”. In: *ArXiv180608792 Astro-Ph Physicsgr-Qc* (Aug. 2, 2018). arXiv: [1806.08792](https://arxiv.org/abs/1806.08792) [astro-ph, physics:gr-qc]. URL: <http://arxiv.org/abs/1806.08792> (visited on 09/14/2021).

- [197] Lei Zhang et al. *Numerical Simulation of Lunar Response to Gravitational Waves and Its 3D Topographic Effect Using the Spectral-Element Method*. Dec. 25, 2025. DOI: [10.48550/arXiv.2512.21667](https://doi.org/10.48550/arXiv.2512.21667). arXiv: [2512.21667](https://arxiv.org/abs/2512.21667) [astro-ph]. URL: <http://arxiv.org/abs/2512.21667> (visited on 01/27/2026). Pre-published.
- [198] Lei Zhang et al. *Super Amplification of Lunar Response to Gravitational Waves Driven by Thick Crust*. Nov. 6, 2025. DOI: [10.48550/arXiv.2511.04067](https://doi.org/10.48550/arXiv.2511.04067). arXiv: [2511.04067](https://arxiv.org/abs/2511.04067) [astro-ph]. URL: <http://arxiv.org/abs/2511.04067> (visited on 01/27/2026). Pre-published.
- [199] Lei Zhang et al. *Thick Lunar Crust Amplifies Gravitational-Wave Signal*. Jan. 23, 2026. DOI: [10.48550/arXiv.2601.16567](https://doi.org/10.48550/arXiv.2601.16567). arXiv: [2601.16567](https://arxiv.org/abs/2601.16567) [gr-qc]. URL: <http://arxiv.org/abs/2601.16567> (visited on 01/27/2026). Pre-published.
- [200] Tianqi Zhao and James M. Lattimer. “Universal Relations for Neutron Star F-Mode and G-Mode Oscillations”. In: *Physical Review D* 106.12 (Dec. 2, 2022), p. 123002. ISSN: 2470-0010, 2470-0029. DOI: [10.1103/PhysRevD.106.123002](https://doi.org/10.1103/PhysRevD.106.123002). arXiv: [2204.03037](https://arxiv.org/abs/2204.03037) [astro-ph.HE]. URL: <http://arxiv.org/abs/2204.03037> (visited on 06/15/2026).