UC Berkeley

UC Berkeley Previously Published Works

Title

Beyond Gaussian assumptions: A new robust statistical framework for gravitational-wave data analysis

Permalink

https://escholarship.org/uc/item/1js4h72h

Journal

Physical Review D, 114(2)

ISSN

2470-0010

Authors

Sasli, Argyro
Karamanis, Minas
Karnesis, Nikolaos
et al.

Publication Date

2026-07-15

DOI

10.1103/n5vd-kzp1

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at https://creativecommons.org/licenses/by/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Beyond Gaussian Assumptions: A new robust statistical framework for gravitational-wave data analysis

Argyro Sasli[1,2,3], Minas Karamanis[4], Nikolaos Karnesis[3], Michael W. Coughlin[1,2], Vuk Mandic[1], Uroš Seljak[4] and Nikolaos Stergioulas[3]
[1] School of Physics and Astronomy, University of Minnesota, 55455 MN, USA
[2] NSF Institute on Accelerated AI Algorithms for Data-Driven Discovery (A3D3), MIT, Cambridge, MA 02139 and University of Minnesota, Minneapolis, MN 55455
[3] Department of Physics, Aristotle University of Thessaloniki, Thessaloniki 54124, Greece
[4] Physics Department, University of California and Lawrence Berkeley National Laboratory Berkeley, Berkeley, CA 94720, USA
(Dated: February 26, 2026)

Many traditional algorithms applied in gravitational-wave astronomy rely on the assumption of Gaussian noise, a condition not always met. To meet this need, this study extends a robust statistical framework, advancing previous work on heavy-tailed likelihoods, that adapts the hyperbolic likelihood method for full frequency domain applications. The framework is designed to maintain high performance under ideal conditions while improving robustness against non-Gaussian noise and outliers in real-world data. We demonstrate the efficacy of this approach through two key case studies. The first case study analyzes a massive black hole binary merger in simulated Laser Interferometer Space Antenna (LISA) data with Gaussian noise, showing that the extended hyperbolic likelihood method performs comparably to the more commonly used Whittle likelihood. The second case study examines a stellar-mass black hole binary merger using real ground-based gravitational-wave data containing non-Gaussian noise or overlapping signals, where our framework exhibits increased robustness and yields more accurate parameter estimations. Our results show that the hyperbolic likelihood better captures the true noise distribution, providing a flexible and physically motivated alternative for GW data analysis across current and future detectors.

arXiv:2602.22074v1 [gr-qc] 25 Feb 2026

I. INTRODUCTION

The detection and analysis of gravitational waves (GWs) has opened a new era in astrophysics, allowing us to probe the universe in unprecedented ways. However, accurate estimation of parameters from GW sources is critically dependent on the quality of the data and the understanding of noise [1–5]. A common assumption in the parameter estimation (PE) process is that the noise in the detector data is Gaussian and stationary. This assumption simplifies the mathematical treatment, making the use of Gaussian likelihood functions a standard practice. However, this practice may introduce significant implications for the accuracy and reliability of the estimated parameters [6, 7], particularly due to the presence of noise and potential outliers in the observed signals.

During the period between November 2019 and March 2020, noise transients, commonly known as "glitches" [8, 9], were observed at an average rate of one every 0.32 minutes at the LIGO Hanford Observatory and one every 1.17 minutes at the LIGO Livingston Observatory. Moreover, 17 out of the 90 gravitational-wave events reported in the first three Gravitational Wave Transient Catalogs occurred within one second of a glitch [10]. Similar glitch rates have also been reported in the O4 run [11]. Importantly, subthreshold glitches (SNR < 5) can still overlap in time and frequency with astrophysical signals, subtly impacting parameter inference despite remaining undetected by conventional glitch-mitigation methods [12]. Such unmitigated noise transients and other artifacts [13–17] pose a significant challenge for both detection pipelines and downstream PE analyses [18–21].

Traditional approaches, such as the widely used Gaussian and/or Whittle likelihood, while effective under ideal conditions, can be sensitive to departures from Gaussian noise. At the same time, noise in GW detectors is not perfectly stationary; its statistical properties can change over time [22]. The Gaussian likelihood assumption fails to capture these temporal variations, potentially leading to underestimates or uncertainties in the PE and a higher probability of systematic errors.

In practice, the data collected by detectors like Advanced LIGO [23], Advanced Virgo [24] and KAGRA [25] often contain various artifacts and non-stationarities, imposing variations in the statistical properties of the noise over time. With the better sensitivities of upcoming detectors, such as ground-based detectors like Einstein Telescope (ET) [26], Cosmic Explorer (CE) [27], NEMO [28] and future spaced-based detectors like LISA [29], TianQin and Taiji [30], signals will overlap in time and frequency, resulting in noise uncertainties [31]. In the case of space-based detectors, calibrating the relevant noise models is not entirely feasible [31] (we refer the reader to [31] and references therein for more details in different data challenges in the case of space-borne detectors), as is done for ground-based detectors [32].

Several analyses have been conducted with the aim of tackling the aforementioned issues [33–46]. Here, we adopt the statistical framework that was previously introduced in [6, 47] and extend it to the whole frequency

domain for PE purposes.

In [6], we introduced a heavy-tailed likelihood, more specifically the Hyperbolic likelihood $\Lambda_{\mathcal{H}}$ [6, 7, 48, 49], derived from the Generalized Hyperbolic distribution [49]. This likelihood offers a robust framework for data analysis, effectively handling data outliers, noise non-stationarities, and potential inaccuracies in the noise Power Spectral Density (PSD) model. We applied this approach to examples from GW astronomy, specifically using synthetic datasets from the adopted LISA mission. This method has been demonstrated to yield more robust results in PE [6, 7], particularly in scenarios where the data deviated from the theoretical estimation of the PSD of the noise. We concluded that the Gaussian likelihood assumption typically underestimates the tails of the noise distribution [6]. This can result in overly optimistic confidence intervals [6], where the true parameter values lie outside the estimated ranges more frequently than predicted. This misinterpretation can hinder the reliability of scientific conclusions drawn from GW observations. However, we limited the application to a small frequency segment and did not expand the analysis to the entire LISA frequency domain.

Subsequently, in [47], we enhanced the statistical framework discussed above by incorporating an analysis using B-splines following a shape-agnostic spectral model to estimate the PSD of the noise, similarly to [50]. Furthermore, we characterized stochastic GW signals [47] by adopting this framework based on a segmented set of hyperbolic parameters. Specifically, our approach leverages the hyperbolic likelihood to investigate the spectral characteristics of the signal and concurrently check for any deviations from Gaussianity. The analysis focuses on a stochastic signal generated by overlapping signatures in LISA band, expected to be cyclo-stationary due to orbital motion. The study recovers non-stationarity as a deviation from Gaussianity, particularly in a specific frequency range [47].

The present study aims to go one step further and extend this robust statistical framework to encompass the full frequency domain and simultaneously characterize the noise properties while performing PE for the given GW source. By broadening the applicability of the heavy-tailed likelihood approach, we seek to enhance the overall robustness and reliability of GW PE across a wider range of scenarios and data characteristics. Among the different types of GW sources, we apply this framework for Binary Black-Hole (BBH) mergers, which is a common source for both space-based [29, 30] and ground-based detectors [23–28], allowing us to test this framework for various GW detectors. Specifically, BBH mergers span a wide range of masses and can be broadly categorized into two groups:

- Supermassive BBH mergers: These involve much larger black holes, with masses ranging from millions to billions of solar masses. They are believed to be the result of galaxy mergers and are a primary target for future space-based detectors [29, 30].
- Stellar-mass BBH mergers: These involve black holes with masses typically ranging from a few to tens of solar masses. They are thought to be the end products of massive star evolution and are primarily detected by ground-based interferometers [23–28].

Detecting and analyzing BBHs allows scientists to probe fundamental questions in physics and astrophysics (see [51] and references therein). The first direct detection of GWs from a stella-mass BBH merger by Advanced LIGO in 2015 (GW150914) [52] opened up an entirely new way of observing the universe. Since then, numerous BBH mergers have been detected [10, 11, 18, 53–58], each contributing to our understanding of these extreme cosmic events (for example, see [59–63]) and further detections will give insight into the nature of our universe [64, 65].

The expected detection rate of GW from BBH is projected to increase significantly in the future as more advanced detectors come online and existing detectors improve their sensitivity. With third-generation detectors [26–28], detection rates could increase to thousands or even millions of events per year, making it very possible to have overlapping low-SNR signals in time and frequency. Similar considerations apply to BNS and BHNS systems, for which the typically longer signal duration in the detector band further increases the probability of signal overlap. For this reason, performing PE of a BBH in a very complex data set could be challenging. Gaussian-like likelihoods are expected to provide biased results, since low-SNR signals would break the Gaussianity of the noise and increase noise uncertainties. We simulate such cases and give more details in Sect. III C.

This paper presents two key case studies that demonstrate the advantages of our framework

1. Analysis of a Massive Black Hole Binary merger (MBHB) in simulated LISA data, characterized by Gaussian noise (see Sect. III B). This case allows us to benchmark the performance of our extended hyperbolic likelihood method against the traditional Whittle likelihood under ideal noise conditions. In addition, in Appendix A, we present a comparison between the Whittle and Gaussian likelihoods.

2. Examination of a stellar-mass black hole binary merger using real ground-based gravitational-wave detector data, containing non-Gaussian noise features (see Sect. III C). This scenario enables us to assess the robustness of our method in handling real-world data complexities and to demonstrate its advantages over conventional approaches.

Through these analyses, we aim to show that this framework maintains high performance under ideal conditions while offering robustness in the presence of non-Gaussian noise and outliers. This advancement has significant implications for improving the accuracy and reliability of gravitational-wave data analysis across various detectors and astrophysical sources.

The paper is structured as follows. Sect. II provides a detailed description of the statistical framework, including the extension of the hyperbolic likelihood method to the full frequency domain. Sect. III outlines our methodology, detailing the data sources, simulation setups, and main results. Sect. IV discusses the implications of our findings for GW astronomy and an outlook on future applications. Appendix A presents a benchmark comparison between Gaussian and Whittle likelihoods in the LISA scenario under controlled noise assumptions, while Appendix B provides supplementary details and results for the multiple-injection and glitch analyses.

II. STATISTICAL FRAMEWORK

In a Bayesian framework, the posterior distribution of the parameters $\vec{\theta}$ is given by

$$p(\vec{\theta}|x) \propto p(x|\vec{\theta})p(\vec{\theta}), \tag{1}$$

where $p(x|\vec{\theta})$ represents the likelihood function and $p(\vec{\theta})$ denotes the prior distribution of the parameters. The marginal likelihood $p(x)$ typically serves as a normalization constant and is therefore omitted from Eq. (1), although it plays a central role in model selection [66].

A. Data model and residuals in GW analysis

In GW data analysis, detector data are modeled as the sum of signal and noise,

$$s(t) = h(t;\vec{\theta}_{\rm source}) + n(t;\vec{\sigma}), \tag{2}$$

where h is the GW signal depending on source parameters $\vec{\theta}_{\rm source}$, and n denotes instrumental noise, which may be described by a set of noise parameters $\vec{\sigma}$.

Given a candidate signal model, the residual data are defined as

$$x(t) = s(t) - h(t;\vec{\theta}_{\rm source}), \tag{3}$$

which represent the noise realization under the assumed signal hypothesis. Likelihoods are constructed from these residuals.

In the GW community, data analysis is commonly performed in the frequency domain, where stationary colored noise is naturally described by the PSD. In this representation, the PSD defines the metric used to measure residual energy through the noise-weighted inner product.

B. Noise-weighted inner product and Gaussian likelihood

The noise-weighted inner product between two real time series a and b is defined as [67, 68]

$$\langle a|b\rangle = 4\,\mathrm{Re}\int_0^\infty \mathrm{d}f\,\left[\tilde{a}^*(f)C_n^{-1}(f)\tilde{b}(f)\right], \tag{4}$$

where $C_n(f)$ is the one-sided noise PSD, the tilde ($\tilde{}$) denotes the Fourier transform, and (*) denotes complex conjugation. For multiple data channels or detectors, C_n becomes a matrix whose off-diagonal components describe cross-spectral noise correlations [6, 47].

The quadratic residual energy appearing in likelihood constructions can then be written as

$$r_i = \mathrm{Re}\left\{\tilde{x}_i^* C_n^{-1}\tilde{x}_i\right\}, \tag{5}$$

which represents the noise-weighted residual power in the frequency domain.

In this framework, the Gaussian log-likelihood is commonly written as

$$\Lambda_{\mathcal{N}}(\vec{\theta}_{\rm source}) \propto -\frac{1}{2}\langle x|x\rangle. \tag{6}$$

The Whittle approximation of the Gaussian likelihood (see, e.g., [6, 43, 69–72]) takes the form

$$\Lambda_{\mathcal{W}}(\vec{\theta}) \propto -\sum_f \left(\ln C(\vec{\theta}) + \langle x|x\rangle\right), \tag{7}$$

with $\vec{\theta} \equiv \{\vec{\theta}_{\rm source}, \vec{\sigma}\}$, where $C(\vec{\theta})$ denotes the parametrized noise covariance and $\vec{\sigma}$ the PSD parameters.

C. Hyperbolic likelihood formulation

To address limitations associated with likelihood models that assume Gaussian noise properties, one may adopt likelihood models with heavier tails. Several approaches have been proposed in the literature, including likelihoods based on Student's t distribution [73, 74], or mixture-based likelihood constructions [44, 45]. Such approaches aim to provide robustness against non-Gaussian noise features and transient deviations from stationarity.

In this work, we adopt the Hyperbolic likelihood model [6, 47–49]. One of the key advantages of this model is its flexibility in describing a wide range of statistical behaviors in the data, including deviations from Gaussianity, while retaining a well-defined likelihood structure. As demonstrated in [6, 47], the $\mathcal{H}$ likelihood is sensitive to departures from Gaussian noise assumptions and can capture heavy-tailed residual statistics. By jointly estimating the parameters of the $\mathcal{H}$ likelihood together with

the physical model parameters, we gain additional information about the statistical properties of the residual data (see Eq. (5)).

For data vectors $\boldsymbol{x}_i \in \mathbb{R}^d$, $1 \leq i \leq n$, where d denotes the dimensionality of the data representation adopted in the likelihood (i.e., the number of real degrees of freedom per sample), the symmetric Hyperbolic log-likelihood is written as [6]

$$\begin{aligned}\Lambda_{\mathcal{H}}(\alpha,\delta) =& n\left[\frac{d+1}{2}\ln\left(\frac{\delta}{\alpha}\right)+\frac{1-d}{2}\ln(2\pi)\right.\\ &\left.-\ln(2\alpha)-\ln\left[K_{(d+1)/2}(\delta\alpha)\right]\right]\\ &-\alpha\sum_{i=1}^{n}\sqrt{\delta^2+r_i},\end{aligned} \tag{8}$$

where

$$r_i \equiv \boldsymbol{x}_i^{\mathrm{T}}\widehat{\Delta}^{-1}\boldsymbol{x}_i, \tag{9}$$

and $\widehat{\Delta}$ denotes an estimate of the symmetric positive definite dispersion matrix defining the quadratic metric of the distribution. In the Gaussian limit, this matrix reduces to the noise covariance, or equivalently the PSD operator in the frequency domain.

The analytic expression for the variance of the symmetric Hyperbolic distribution is given by [6]

$$\sigma^2 = \frac{\delta K_2(\delta\,|\alpha|)}{\alpha K_1(\delta\,|\alpha|)}, \tag{10}$$

where K_1 and K_2 are modified Bessel functions of the third kind. In the Gaussian limit, $\delta, \alpha \to \infty$ such that $\delta/\alpha \to \sigma^2$. In the univariate time-domain formulation, σ^2 reduces to the noise variance, whereas in the frequency-domain GW formulation this generalizes to the frequency-dependent noise PSD [6, 7].

In gravitational-wave applications, precisely defining the dimensionality d is important. For n_c independent data channels (detectors), we take $d = n_c$ in the time domain. In the frequency domain, each channel contributes one complex Fourier coefficient, corresponding to two real degrees of freedom, and therefore we set $d = 2n_c$ per frequency bin [47].

Depending on the frequency range considered in the analysis, we may adopt a piecewise parametrization of the Hyperbolic model in order to achieve a robust reconstruction of the effective noise PSD. Specifically, the frequency domain can be divided into N_s frequency bins, within which the Hyperbolic parameters are inferred independently. Therefore, the parameters can be taken to be $\alpha_{N_s}, \delta_{N_s}$, α, δ_{N_s}, or α_{N_s}, δ. For computational efficiency, in this work we adopt the α, δ_{N_s} parametrization.

In conventional GW parameter estimation, one may introduce explicit noise model parameters $\vec{\sigma}$ to reconstruct the PSD of the noise, leading to an inferred parameter set $\vec{\theta} \equiv \{\vec{\theta}_{\mathrm{source}}, \vec{\sigma}\}$. In contrast, when adopting the $\mathcal{H}$ likelihood, only minimal assumptions about the detailed noise model are required, and inference is instead performed directly on the parameters of the $\mathcal{H}$ distribution. In this way, the Hyperbolic parameters provide an effective statistical description of the residual noise properties while simultaneously enabling robust parameter estimation of the physical source model.

III. ANALYSIS

Our previous studies demonstrated the robustness of the hyperbolic likelihood in handling both Gaussian and non-Gaussian noise. For Gaussian noise with a single data source, both the $\mathcal{H}$ and $\mathcal{W}$ likelihoods performed similarly regardless of the noise modeling accuracy, while the Gaussian likelihood was highly dependent on the correct PSD of the noise [6]. In non-Gaussian noise conditions, the $\mathcal{H}$ likelihood accurately recovered the noise variance, in contrast to the $\mathcal{W}$ likelihood, which overestimated it [47] due to its Gaussian noise assumption. In addition, the importance of extracting the underlying statistical properties of the data is highlighted in [47]. For example, the detection and the characterization of non-Gaussian stochastic signals from ultra-compact binaries could hinder the underlying physical processes and the population of these sources in the universe [47].

For the purposes of this paper, we adopt a combined methodology from our previous studies. We perform PE for BBH under both ideal (Sect. III B) and non-ideal (Sect. III C) noise conditions, following the statistical framework described in Sect. II, while accounting for corrections in the assumed PSD noise. We assume that the data channels are statistically uncorrelated, such that the total log-likelihood factorizes into the sum of the log-likelihoods of the individual channels. This approximation is well justified for the datasets considered here. For LISA, we analyze the approximately independent A and E TDI channels [75, 76], while for ground-based observations we treat data from independent detectors, such as LIGO Livingston (L1) and Hanford (H1), as uncorrelated.

The PE with hyperbolic likelihood $\mathcal{H}$ is compared with Whittle $\mathcal{W}$ and Gaussian likelihood. For the $\mathcal{W}$ likelihood, we assume a simple noise model; a level correction on the assumed PSD of the noise, $\vec{\sigma}$. An alternative case could be the assumption of the actual LISA noise model, and infer the nominal values for the SciRD noise, for example see Appendix IV A. The dimensionality of this parameter depends on the number of frequency bins N_s we divide into the frequency domain. Therefore, the noise parameters correspond to band-dependent PSD correction factors. For the $\mathcal{H}$ likelihood, we introduce band-dependent δ_{N_s} parameters, while assuming a common α parameter across the full frequency band, consistent with the formulation described in Sect. II. This results in N_s additional noise parameters for the $\mathcal{W}$ likelihood and $N_s + 1$ for the $\mathcal{H}$ likelihood.

A. Sampling strategy

For PE purposes, we employ two packages, `Eryn` [77] and `pocoMC` [78]. `Eryn` utilizes a Parallel Tempering Markov Chain Monte Carlo (PTMCMC) approach [77, 79], which is particularly effective for gravitational wave inference problems where posterior distributions may exhibit strong correlations, non-Gaussian structures, and multi-modality. By simultaneously evolving multiple ensembles of Markov chains at different inverse temperatures, `Eryn` facilitates efficient exploration of complex likelihood landscapes. The exchange of chains at neighboring temperatures helps low-temperature chains escape local maxima, thereby improving convergence and robustness. Initialized with walkers drawn from small perturbations around the injected parameters, `Eryn` accelerates burn-in while ensuring thorough exploration of the posterior. Proposals are generated using affine-invariant ensemble moves, with samples from the lowest-temperature chain representing the target posterior distribution, providing a reliable baseline for posterior sampling.

`pocoMC` is a Preconditioned Monte Carlo sampler [78] that combines annealed Persistent Sampling [80] with normalizing flows, which are trained on-the-fly to learn an approximate transport map from the target distribution to a latent space with reduced correlations. This method significantly enhances sampling efficiency compared to traditional MCMC or nested sampling methods. The algorithm evolves a population of particles through a sequence of intermediate distributions, updating the normalizing flow to effectively "Gaussianize" the distribution and enable large, informed proposal steps. `pocoMC` has been validated on gravitational wave inference problems using `bilby`, demonstrating substantial speedups and accurate posterior reconstruction [81, 82]. In this work, we further improve `pocoMC`'s performance with a fully vectorized likelihood implementation, making it particularly well-suited for computationally demanding analyses with complex likelihoods.

B. LISA scenario: Simulated Gaussian noise

We investigate a low-SNR case from the LISA Data Challenge (LDC) *Sangria* data set [83], corresponding to a massive black hole binary (MBHB) signal. Waveforms are generated using the `BBHx` package [84–86] with the `PhenomD` approximant [87]. The injected waveform parameters are listed in Table I. The simulated data have a total observation duration of $T_{\rm obs} = 30.4368$ days, while the coalescence time is set to one day before the end of the measurement.

The inferred source parameters $\vec{\theta}_{\rm source}$ correspond to those listed in Table I. However, instead of sampling directly in the component masses (m_1, m_2), we sample in the chirp mass $\mathcal{M}$ and mass ratio q $(\equiv m_1/m_2)$.

Considering the LISA constellation and instrument, we assume a rigid constellation, and equal noise across all space-crafts. This allows us to use the noise-orthogonal set of TDI channels (A, E and T) [88, 89], computed from the X, Y, Z channels as

$$A = \frac{1}{\sqrt{2}}(Z-X), \quad E = \frac{1}{\sqrt{6}}(X-2Y+Z), \quad T = \frac{1}{\sqrt{3}}(X+Y+Z). \tag{11}$$

The analytic expression for the average noise level in the A, E TDI channel is

$$PSD^{\rm AE}(f) = 8\sin(\hat{\omega})^2\left\{2S_{\rm i}\left[3+2\cos(\hat{\omega})^2+\cos(2\hat{\omega})\right] + S_{\rm a}\left[2+\cos(\hat{\omega})\right]\right\}, \tag{12}$$

where $\hat{\omega} = 2\pi f L/c$, f the frequency and L is the average detector arm-length. The noise components S_i and S_a are fixed for the general case and selected to be in agreement with the nominal values for the SciRD noise [90, 91],

$$S_a = \log_{10}\left(3\cdot 10^{-15}\right)^2 \quad \frac{\rm m^2}{\rm Hz\cdot sec^4}, \qquad S_i = \log_{10}\left(15\cdot 10^{-12}\right)^2 \quad \frac{\rm m^2}{\rm Hz}, \tag{13}$$

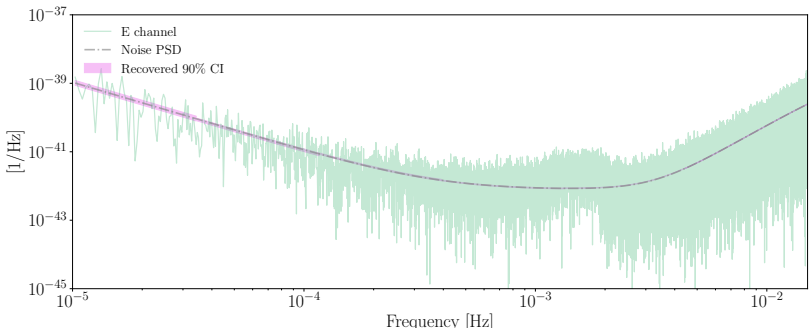
which are expressed in relative frequency units [76]. For this simple investigation, the T channel is neglected since it greatly suppresses the GW signal.

Parameters	SNR = 357
Mass 1, m_1 $[M_\odot]$	4956676.2876
Mass 2, m_2 $[M_\odot]$	4067166.60352
Spin of body 1, χ_1	-0.523732715
Spin of body 2, χ_2	-0.117412144
Distance, D [Mpc]	61097.116076
Inclination, ι $[rad]$	1.420048341
Ecliptic Latitude, β $[rad]$	-1.081082148
Ecliptic Longitude, λ $[rad]$	4.052962883
Polarization, ψ	1.22844350008
Phase at coalescence, ϕ_0	0.6417162631

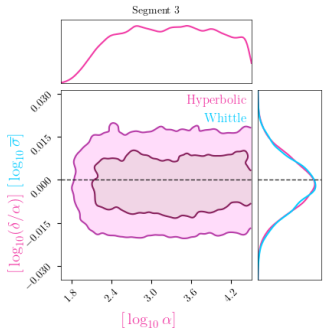
TABLE I: Parameters of our model from Sangria training data-set. The SNR of this case is 357 for a one-year data.

$\vec{\theta}_{\rm source}$	**Lower bound**	**Upper bound**
$\mathcal{M}$	$0.39\cdot 10^6$	$7.8\cdot 10^6$
q	0.99999	2
χ_1	-1	1
χ_2	-1	1
$\log_{10}$ D	$0.1\log_{10}$ $\rm D_{true}$	$2\log_{10}$ $\rm D_{true}$
$\rm t_c$	365.2422	8765.8128
$\cos\iota$	-1	1
$\sin\beta$	-1	1
λ	0	2π
ψ	0	π
ϕ_0	0	2π

TABLE II: The prior settings for the parameters $\vec{\theta}_{\rm source}$.



(a) Recovered PSD in the LISA band using the $\mathcal{H}$ likelihood.



(b) Example of PSD correction for a randomly selected frequency segment. Segment 3 corresponds to the third frequency bin used in the analysis and is shown with the $\mathcal{H}$ (pink) and $\mathcal{W}$ (blue) likelihoods. Both predict no correction to the adopted PSD.

FIG. 1: PSD analysis in the LISA band.

We run the PTMCMC with 20 temperatures, 80 walkers, and 3×10^5 samples per walker, while we set the burn-in samples equal to 1×10^5. The inferred parameters are the source parameters and the noise fitting parameters for N_s=6 equally in log spaced frequency bins. This results in a total parameter space $\vec{\theta}$ of 19 when account for $\mathcal{W}$ likelihood, and 20 for $\mathcal{H}$ likelihood.

Finally, we have adopted wide priors for the source parameters, which are summarized in Table II. The parameters with which we perform the noise fitting as well as the first eight source parameters follow the uniform distribution, and the remaining follow a periodic prior. The whole PE is performed on GPUs, using cuda [92], increasing the efficiency up to 10 times compared to CPU-based analyses.

Figure 1a illustrates the PSD reconstruction for the LISA scenario. The plot shows the true noise PSD (black line) along with the recovered 90% confidence interval (shaded purple region) obtained using the $\mathcal{H}$ likelihood. Since both likelihoods converge to the Gaussian limit in this regime, the reconstructed PSDs from the $\mathcal{H}$ and $\mathcal{W}$ likelihoods are statistically indistinguishable, and we therefore show only the $\mathcal{H}$ reconstruction for clarity. Both likelihood methods demonstrate excellent performance in reconstructing the PSD across the frequency range of interest (10^{-5} to 10^{-2} Hz). The recovered PSD closely follows the true noise PSD, with the 90% confidence interval encompassing the true PSD line over most of the frequency range.

In Fig. 1b, we show the PSD correction posteriors for a representative frequency segment. The posteriors obtained using both likelihoods indicate that no significant refinement of the PSD is required in this segment. The distributions for both methods (pink for $\mathcal{H}$ and blue for $\mathcal{W}$) are relatively compact and centered near the true value, corresponding to zero logarithmic PSD correction.

Parameters	**$\mathbf{SNR_{L1}}$ = 24** **$\mathbf{SNR_{H1}}$ = 14**
Mass 1, m_1 [$M_\odot$]	36
Mass 2, m_2 [$M_\odot$]	29
Distance, D [Mpc]	1000
Right Ascension (RA)	1.375 rad
Declination (Dec)	0.2108 rad
Polarization, ψ	2.659 rad
Phase at coalescence, ϕ_0	1.3 rad
Inclination angle, [rad]	0.4
Primary object tilt [rad]	0
Secondary object tilt [rad]	0
Primary dimensionless spin magnitude	0
Secondary dimensionless spin magnitude	0
Relative spin azimuthal angle, ϕ_{12} [rad]	0
Spin phase angle, ϕ_{jl} [rad]	0

TABLE III: Parameters of our targeted BBH injection in LIGO data.

We further compared the performance of the $\mathcal{H}$ likelihood method with the traditional $\mathcal{W}$ likelihood. The results demonstrate that both methods perform comparably well in this scenario, characterized by Gaussian noise. Fig. 2 presents the posterior distributions for various parameters of the binary system. The overlapping posterior distributions for the $\mathcal{H}$ and $\mathcal{W}$ likelihoods indicate that both methods provide consistent PEs. This similarity in performance suggests that the proposed $\mathcal{H}$-likelihood analysis converges to a likelihood shape that closely resembles the Gaussian, thus the similar results to the whtl likelihood analysis.

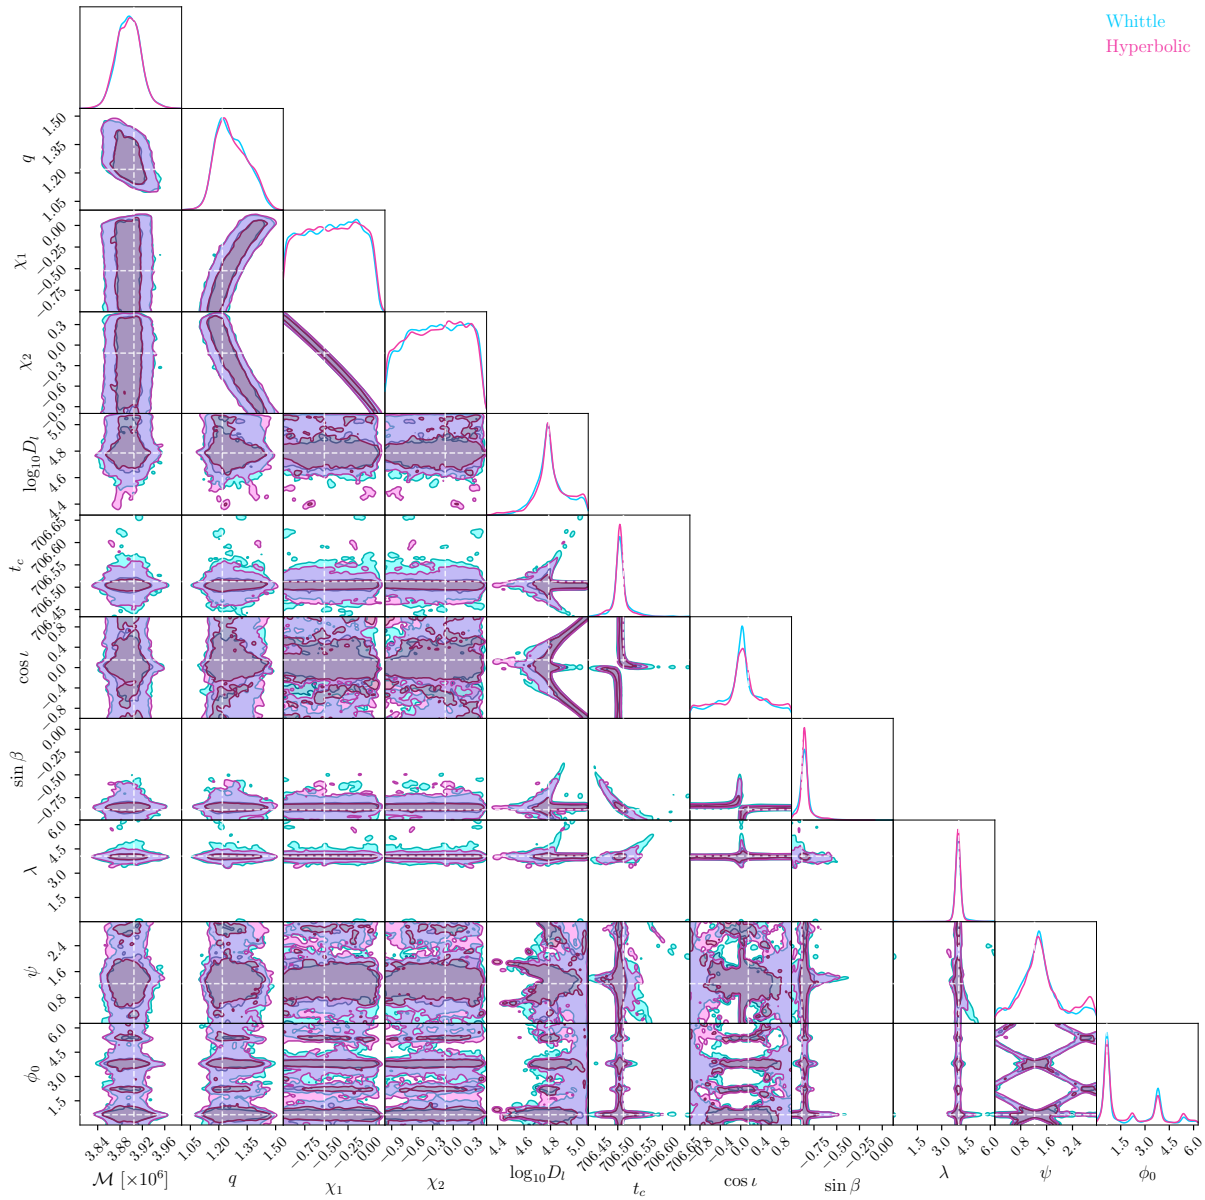


FIG. 2: Cornerplots using the $\mathcal{W}$ (blue) and the $\mathcal{H}$ (pink) likelihood. This case correspond to a source of 357 SNR for one-year data (see Table I). Units follow Table I.

C. Ground-based detector scenario: Real noise

We investigate three different cases of real noise, but with more complex scenarios. These cases consist of datasets where we have a targeted BBH in the presence of (1) seven sub-threshold SNR BBHs, (2) a long-lasting low-SNR BBH, and (3) a noise transient (glitch). For each dataset, the PSD is estimated from 128 s of data using a Welch-style approach with 8 s Fourier segments windowed with a Tukey window ($\alpha = 0.2$) and 50% overlap, yielding 31 segments. Individual periodograms are computed for each segment and combined using the median at each frequency bin to obtain a robust PSD estimate.

For BBH waveform modeling and signal injections into real data, we employed the `bilby` package [93] and adopted `IMRPhenomPv2` [94]. This allows us to simulate gravitational-wave signals within the context of actual detector noise, providing a realistic scenario for testing our analysis methods. For every case, we inject the same signal, a simulated BBH with parameters as shown in Table III. Those parameters were chosen to represent a typical stellar-mass BBH system detectable by ground-based detectors.

When running with `Eryn`, we use 20 temperatures, 50 walkers, and 7×10^4 samples per walker, with 5×10^4 burn-in samples. For `pocoMC`, we set $n_{\rm eff} = 2000$,

$n_{\text{active}} = 1000$, and use 5×10^4 total samples, corresponding to a conservative configuration. In this setup, `pocoMC` achieves approximately a fourfold speed-up compared to `Eryn`.

All parameter estimation runs are performed on CPUs, with waveform and likelihood computations parallelized across approximately 70% of the available cores (32 CPUs), resulting in an effective speed-up of up to a factor of three.

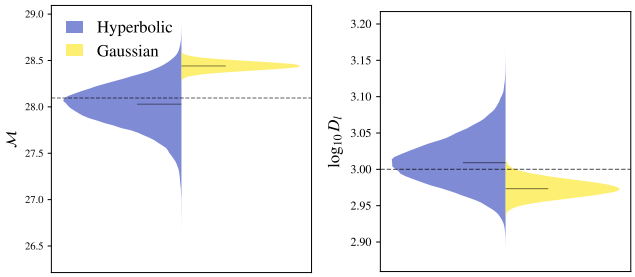
1. *Multiple sources of BBH*

We selected a segment of Advanced LIGO data centered on GPS time `1238189526.951953`, which is close to a glitch and thus the data contain non-stationary noise when the PSD is calculated. We then investigate two different cases, that could make our PE challenging. The cases are as follows:

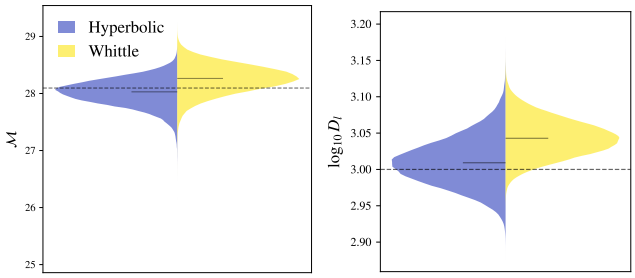
- *Injecting a long lasting BBH source:* We generate a sub-threshold BBH signal (see parameters in Table VI) with a duration of 5.7 seconds, which exceeds the 4-second data segment used in our analysis. The long-duration BBH is injected 0.09 seconds prior to the trigger time of the "target" BBH event. This configuration poses a challenging scenario for likelihood formulations that do not account for PSD uncertainty, since the extended signal contaminates the data segment and effectively alters the stationary noise and background properties.

- *Injecting seven different low SNR BBH sources*: We performed injections of seven low-SNR BBH signals with a common trigger time of `1238189542.951953`, and the waveform parameters were randomly generated (and Table VII). The choice of a common trigger time is intentional. By aligning all sub-threshold BBH signals with the trigger time of the targeted BBH event, we construct a controlled scenario in which multiple weak signals coherently overlap within the same data segment. This setup effectively produces an artificial increase in the apparent signal content at the trigger time, creating a challenging environment for likelihood models. In particular, it allows us to probe the robustness of different likelihood formulations under conditions where the background noise is contaminated by unresolved overlapping signals.

Figures 3 and 4 demonstrate that the $\mathcal{H}$ likelihood reduces the bias for the key binary parameters, chirp mass $\mathcal{M}$ $[M_\odot]$ and luminosity distance. Posterior distributions for these cases and waveform reconstructions in the frequency domain are presented in Appendix B with Figures 10, 11 and 12.



(a) Hyperbolic (blue) and Gaussian (yellow) likelihood



(b) Hyperbolic (blue) and Whittle (yellow) likelihood

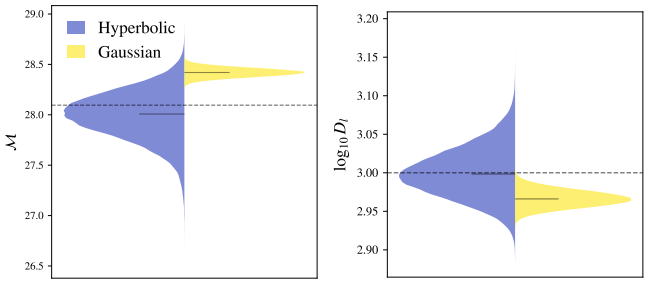
FIG. 3: Posterior distributions obtained with the Hyperbolic compared to (a) Gaussian and (b) Whittle likelihoods. Each panel shows split violin plots for the most affected key binary parameters, namely chirp mass $\mathcal{M}$ $[M_\odot]$ and luminosity distance. This case refers to the data with the long-lasting BBH injected. The hyperbolic likelihood systematically reduces bias.

2. *Glitches*

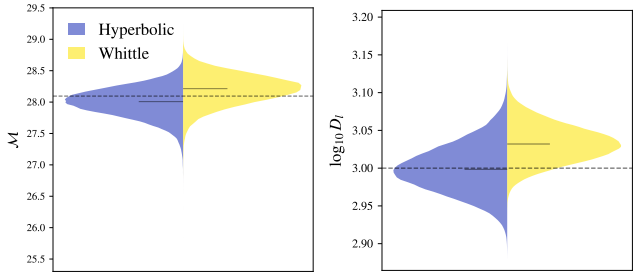
In Ref. [19], both successful and unsuccessful cases of PE in the presence of glitches were demonstrated using the `BayesWave` pipeline. Here, we deliberately focus on a representative case in which the PE remains biased despite the application of glitch-mitigation techniques. Specifically, we consider the "blip" glitch with trigger time `1165578732.45`, where a BBH signal is injected 0.025 s prior to the glitch onset. For this configuration, the method presented in Ref. [19] fails to recover the injected parameters: the true values are excluded at the 90% credible level in both the pre-deglitching and post-deglitching posterior distributions. Fig. 5 shows the Q-plot for this data segment from H1 channel (LHO detector) with the injected BBH from Tab. III.

Fig. 6 illustrates the frequency domain reconstruction. In this case, only the $\mathcal{H}$ likelihood successfully recovers the injected BBH signal, with the posterior median closely tracking the true waveform across the full frequency range and the injected signal remaining well within the 90% C.I.

The Gaussian likelihood fails entirely in the presence of the glitch, yielding a reconstruction that is strongly biased and does not overlap with the injected signal. The $\mathcal{W}$ likelihood provides a reconstruction that is qualita-



(a) Hyperbolic (blue) and Gaussian (yellow) likelihood



(b) Hyperbolic (blue) and Whittle (yellow) likelihood

FIG. 4: Similar to Fig. 3, but for the case with the seven low-SNR BBH injections..

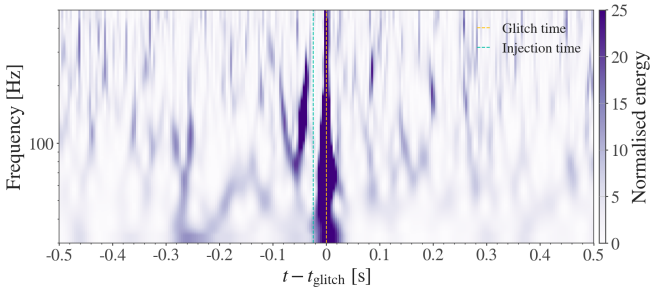
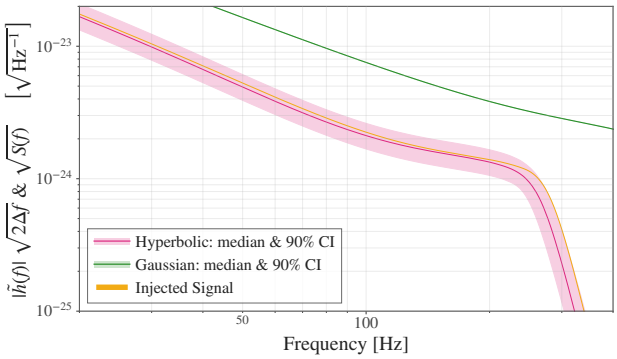
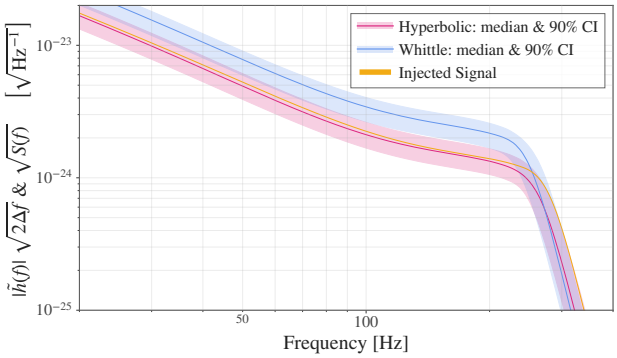


FIG. 5: Q-plot for `1165578732.45` trigger time with the BBH injection. The red dashed line correspond to the glitch time and the green one to the injected signal.



(a) Hyperbolic (pink) and Gaussian (green) likelihood



(b) Hyperbolic (pink) and Whittle (blue) likelihood

FIG. 6: Similar to Fig. 11, but for the case with the glitchy data. Only the hyperbolic likelihood provides accurate coverage of the injected signal, while the Gaussian and Whittle likelihoods fail to enclose the signal within their respective 90% credible intervals.

tively closer to the true signal; however, the injected waveform lies outside the corresponding 90% C.I. over a substantial fraction of the frequency band, indicating incorrect uncertainty quantification.

Fig. 7 shows the time-domain reconstruction of the injected BBH signal recovered using the $\mathcal{H}$ and $\mathcal{W}$ likelihoods. In the presence of the glitch, the hyperbolic likelihood yields a reconstruction which is fully encapsulated within the 90% C.I. In contrast, the $\mathcal{W}$ likelihood produces a reconstruction that is visibly biased in phase and amplitude, with the injected signal lying outside the corresponding C.I. over a significant portion of the time series.

Fig. 8 shows the violin plots for some key recovered parameters, useful either for population studies (chirp mass, mass ratio spins) or sky-location for follow-up strategies. The hyperbolic likelihood systematically reduces bias and posterior broadening compared to the Whittle approximation.

IV. CONCLUSION & DISCUSSION

In our previous work [6], we introduced a heavy-tailed likelihood based on the Generalized Hyperbolic (GH) distribution, known as hyperbolic likelihood $\Lambda_{\mathcal{H}}$. This framework was shown to provide robust parameter estimation in the presence of noise outliers, non-stationarities, and PSD model inaccuracies. Using LISA-like synthetic datasets, we demonstrated that the standard Gaussian likelihood underestimates the true noise tails, leading to over-confident posterior intervals and reduced reliability of inferred astrophysical parameters.

We subsequently extended this approach in [47] by incorporating a shape-agnostic PSD modeling technique based on B-splines, following a strategy similar to [50]. Within this unified heavy-tailed framework, we further analyzed stochastic GW signals in the LISA band and showed that deviations from stationarity and Gaussianity can be consistently recovered through the hyperbolic

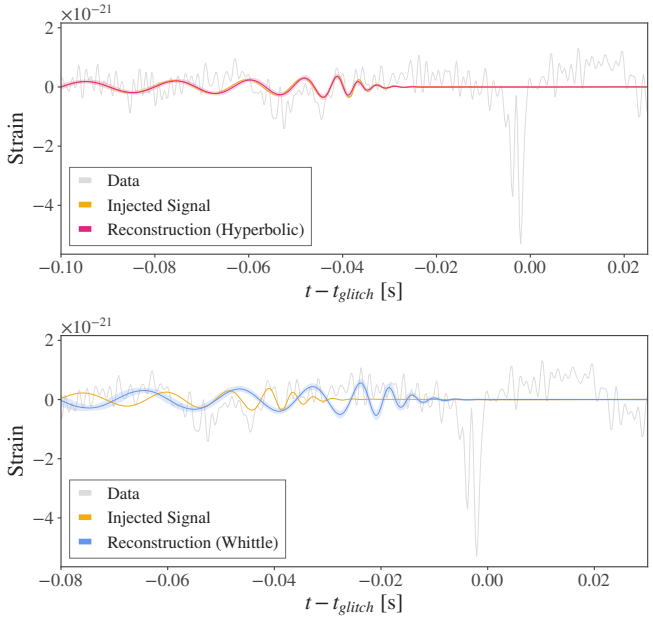


FIG. 7: Time-domain reconstructions of the targeted BBH signal using the $\mathcal{H}$ likelihood (top, pink) and the $\mathcal{W}$ likelihood (bottom, blue). In each panel, the solid line shows the posterior median and the shaded region denotes the 90% credible interval. The injected signal is shown in yellow, while the gray trace corresponds to the strain data. The hyperbolic likelihood provides a reconstruction that closely follows the injected waveform and yields reliable uncertainty coverage, whereas the Whittle likelihood fails to recover the parameters of the injected signal within its inferred CI.

likelihood formulation and its associated spectral parameters.

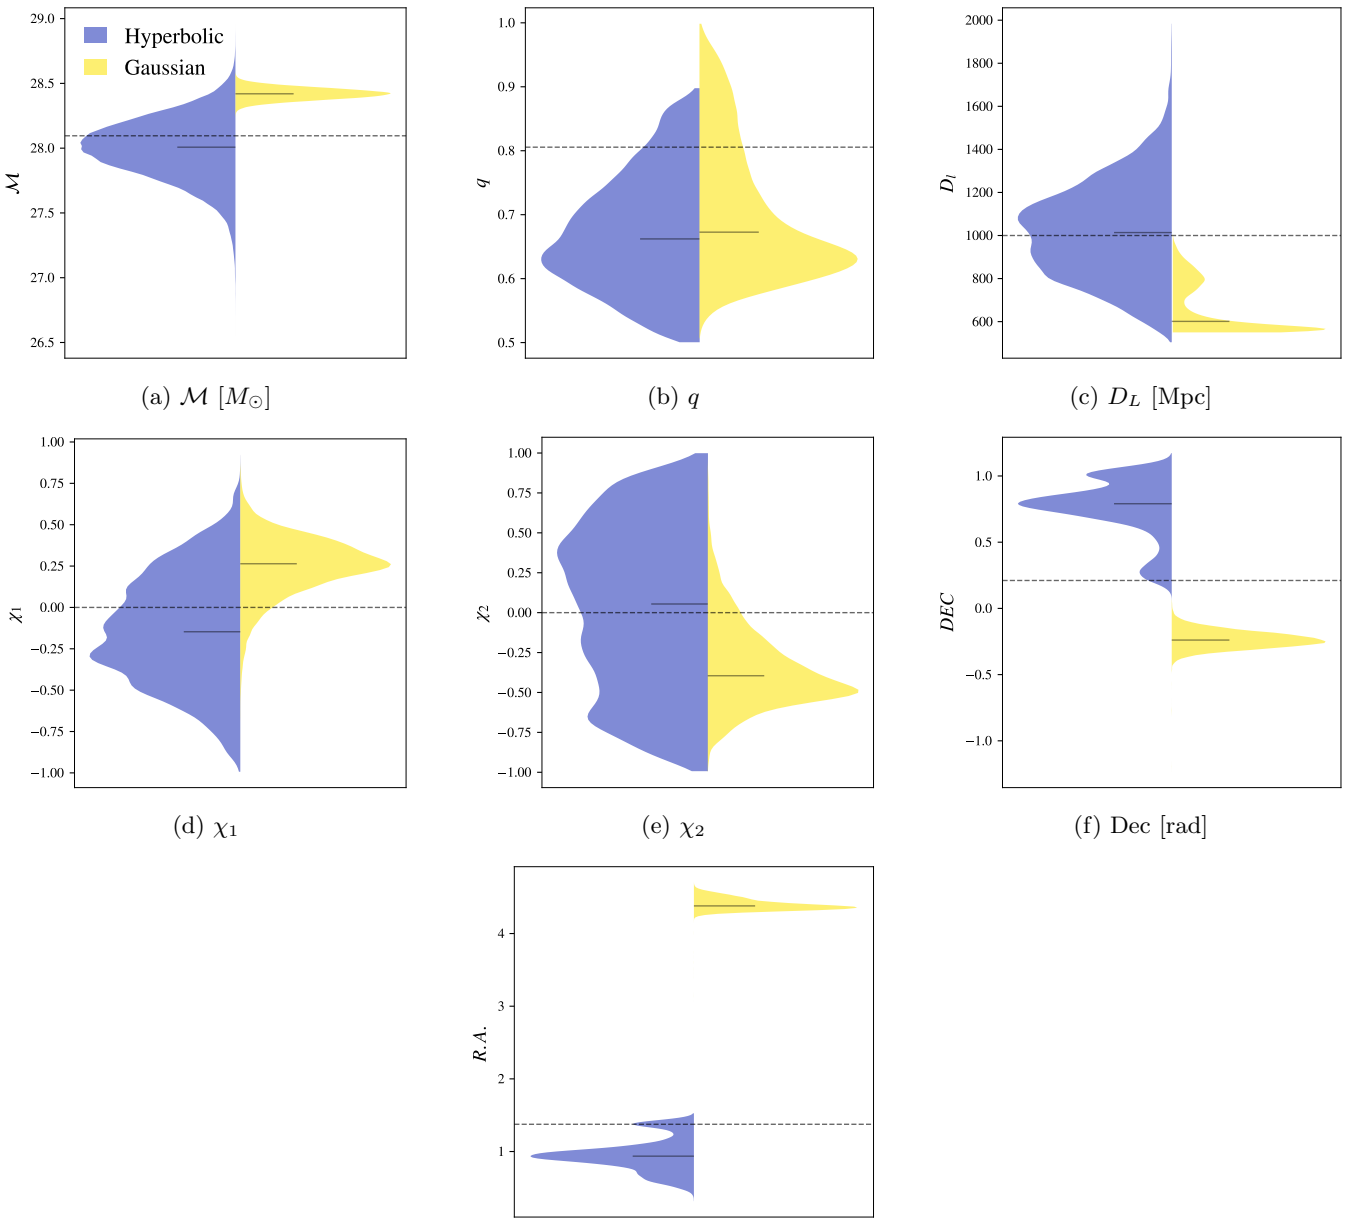
Building on these developments, the present work generalizes the hyperbolic likelihood framework to the full frequency domain, providing a practical and scalable solution for robust GW parameter estimation. This extension allows us to relax the commonly adopted Gaussian noise assumption while retaining high performance under ideal conditions and significantly improving robustness in the presence of data imperfections. Using simulated LISA data with stationary Gaussian noise, we show that the $\mathcal{H}$ likelihood method performs the same as the $\mathcal{W}$ likelihood under Gaussian noise conditions. Both methods provide identical parameter estimation results and demonstrate capability in reconstructing the noise PSD.

We further apply this framework to real ground-based detector data. Our focus is on GW data featuring overlapping signals in time and frequency, as well as non-Gaussian noise transients and glitches, conditions that are expected to become ubiquitous in next-generation ground-based detectors such as Cosmic Explorer and in space-based observatories like LISA. Our investigation confirms that the Gaussian likelihood fails to recover the injected signals in these non-ideal conditions, producing biased reconstructions and severely underestimated uncertainties. The Whittle, $\mathcal{W}$, likelihood yields qualitatively improved reconstructions relative to the Gaussian case, but still suffers from incorrect coverage, when the dataset contains fast noise transients. Importantly, the hyperbolic, $\mathcal{H}$, likelihood achieves this robustness without requiring an explicit parametric model for the noise PSD. Instead, deviations from Gaussianity are captured directly through the inferred hyperbolic parameters, providing a more robust PE. This feature is particularly relevant for future detectors, where non-stationary noise, calibration uncertainties, and signal confusion are expected to be the norm rather than the exception. A corner plot comparing the $\mathcal{H}$ and $\mathcal{W}$ likelihood is given in Appendix B (13).

As an alternative to modifying the likelihood function, one could transform the data in order to Gaussianize them prior to analysis. An extreme case of this methodology involves the complete excision of non-Gaussian segments, such as data containing noise artifacts like glitches. More generally, it involves applying a mathematical transformation to the data before they are used, effectively enforcing a Gaussian distribution. Such data transformation strategies have been successfully applied in other domains, such as the vetting of exoplanet transits (e.g., see [95], and Appendix F of [96]). However, transforming the data in a pre-processing step assumes a fixed transformation, thereby ignoring any uncertainty regarding what constitutes the optimal adjustment. This is conceptually akin to utilizing a generalized likelihood function—such as the hyperbolic likelihood we employ—but artificially fixing its parameters instead of sampling them. By jointly sampling the parameters of the hyperbolic function, our approach natively accounts for the uncertainty in the noise distribution, effectively capturing the uncertainty in the Gaussianization of the data. A rigorous quantitative comparison of these methodologies is deferred to future investigations.

To ensure computational scalability, we implemented a vectorized version of the hyperbolic likelihood evaluation. This enables parallel evaluation of the likelihood over large ensembles of parameter samples, thereby substantially reducing the computational cost of Bayesian inference in high-dimensional parameter spaces. This computational optimization facilitates the integration of the proposed framework into end-to-end inference pipelines, which will become available to the public in the near future.

The package is currently under active development and will be made publicly available on GitHub [97] upon publication of this work. The initial release will include the core likelihood implementation, PSD estimation utilities, waveform generation on both CPU and GPU architectures, and example parameter estimation workflows for real data applications. Future development plans include benchmarking against standard GW parameter estimation pipelines, extending the framework to stochastic GW searches (using LIGO data), and ensuring full



(a) $\mathcal{M}$ $[M_\odot]$ (b) q (c) D_L [Mpc]

(d) χ_1 (e) χ_2 (f) Dec [rad]

(g) RA [rad]

FIG. 8: Comparison of posterior distributions obtained with the hyperbolic and Whittle likelihoods. Each panel shows split violin plots for key binary parameters. The hyperbolic likelihood systematically reduces bias and posterior broadening compared to the Whittle approximation.

compatibility with LISA and other upcoming detectors.

ACKNOWLEDGMENTS

We wish to thank M. Katz, Q. Baghi, and Jean-Baptiste Bayle for useful comments and fruitful discussions. We thank Xiao-Xiao Kou for useful comments on the manuscript. AS acknowledges the Bodossaki Foundation for support in the form of a PhD scholarship. AS and MWC acknowledge support from the National Science Foundation with grant numbers PHY-2308862 and PHY-2117997. NK acknowledges the funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 101065596. N. S. and N. K. acknowledge support from the Gr-PRODEX 2019 funding program (PEA 4000132310). N. S. acknowledges funding from the H.F.R.I. Project No. 26254. In addition, this publication is part of a project that has received funding from the

European Union's Horizon Europe Research and Innovation Programme under Grant Agreement No 101131928. MK and US acknowledge funding from NSF Award Number 2311559, and from the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Contract No. DE-AC02-05CH11231 at Lawrence Berkeley National Laboratory to enable research for Data-intensive Machine Learning and Analysis. We are grateful for the computational resources provided by the LIGO Laboratory and supported by the U.S. National Science Foundation Awards PHY-0757058 and PHY-0823459. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan.

APPENDICES

A. Gaussian vs. Whittle in the LISA band

In this section, we perform PE considering two scenarios. The first concerns the estimation of the MBHB parameters (I) for which we have assumed that the levels of instrumental noise are correct and fully known, while the second refers to the estimation of two additional parameters for the instrumental noise.

The methodology is similar to Sect.(III B), but we perform the analysis by adopting two different likelihoods; the Gaussian and the Whittle $\Lambda_{\mathcal{W}}$ (see Eq. 7). In addition, the noise model differs from the assumed in Sect.(III B), here we infer the actual parameters of the SciRD noise of Eqs.(12,13). For these parameters, we adopt uniform priors, as $\log_{10} S_a \sim \mathcal{U}\{-32,-24\}$ and $\log_{10} S_i \sim \mathcal{U}\{-25,-19\}$.

The results of the analysis are summarized in Fig. 9, and in table V. Considering the above, we can safely conclude that the analysis for both scenarios (of known and unknown noise) yields similar estimations for the parameters. To verify this, we have also computed the Jensen-Shannon Divergence between the posteriors of each case, which are summarized in table IV. The Jensen-Shannon divergence is a measure of how similarly distributed two-different posteriors are [99]. Values close to zero indicate that the distributions do not diverge, while positive values show that they diverge. We find no statistically significant differences between the posterior distributions obtained under the two likelihood assumptions.

$\vec{\theta}$	**Jensen-Shannon Divergence** $[10^{-3}]$
$\mathcal{M}$ $[M_\odot]$	1.24
q	0.89
χ_1	1.01
χ_2	1.03
$\log_{10} D_l$	9.81
t_c	17.83
$\cos\iota$	3.15
$\sin\beta$	13.13
λ	24.77
ψ	10.48
ϕ_0	9.23

TABLE IV: Jensen-Shannon Divergence between the two likelihood types, i.e. the scenario of known (Gaussian likelihood), and unknown noise levels, Eq.(7), for the case in table I. Values are reported in units of 10^{-3}.

$\vec{\theta}_{1,2}$	**Known noise**	**Fitting for the noise**
$\mathcal{M}$ $[M_\odot]$	$\left(388.2^{+2.3}_{-1.1}\right)\times10^4$	$\left(388.3^{+2.1}_{-1.3}\right)\times10^4$
q	$1.228^{+0.107}_{-0.058}$	$1.268^{+0.073}_{-0.095}$
χ_1	-0.3184	$-0.33^{+0.39}_{-0.29}$
χ_2	$-0.41^{+0.59}_{-0.32}$	$-0.34^{+0.48}_{-0.45}$
$\log_{10} D_l$	$4.78^{+0.16}_{-0.12}$	$4.779^{+0.150}_{-0.082}$
t_c [hours]	$706.504^{+0.021}_{-0.024}$	706.504 ± 0.015
$\cos\iota$	$0.01^{+0.41}_{-0.33}$	0.0099
$\sin\beta$	-0.855	$-0.851^{+0.056}_{-0.069}$
λ [rad]	$4.01^{+0.18}_{-0.21}$	$4.01^{+0.15}_{-0.16}$
$\log_{10} S_{\rm a}$	–	$\left(-29038.1^{+5.6}_{-6.5}\right)\times10^{-3}$
$\log_{10} S_{\rm i}$	–	$\left(-21646.6^{+2.6}_{-3.2}\right)\times10^{-3}$

TABLE V: The PE results correspond to the two different log-likelihood types for known and unknown noise levels. The left "Known noise" column refers to the case where the classic Gaussian likelihood is adopted, whereas the right part of the table, named "Fitting for the noise", refers to the analysis using the log-likelihood of Eq.(7).

B. Supplementary details of the LIGO cases

This appendix presents supplementary material supporting the multiple-injection analyses discussed in Sect. III C 1. We include additional corner plots and tables of injected parameters that are not shown in the main text in order to maintain clarity and readability. These results demonstrate the consistency of the inferred posteriors across different injection realizations and provide full transparency regarding the simulation setup.

For the multiple low-SNR injection study, the parameters of the additional BBH signals were drawn randomly from astrophysically motivated prior distributions. The injected sources span a range of masses, sky locations, and orientations, leading to individual signal-to-noise ratios below the typical detection threshold. The full list of injected parameters is reported in Tables III and VII

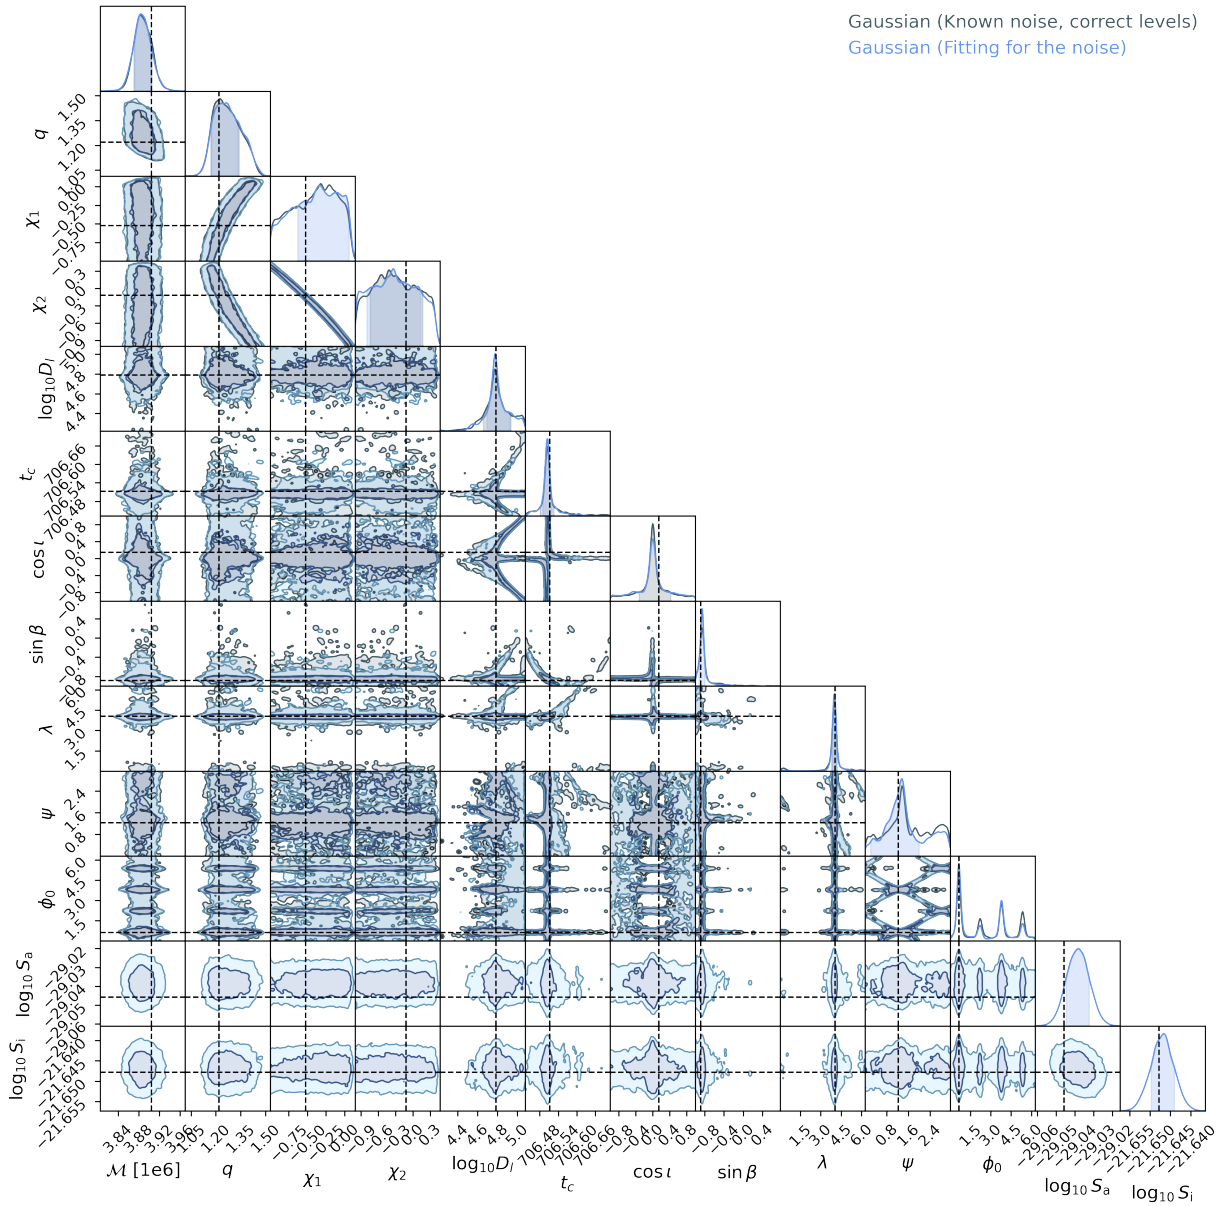


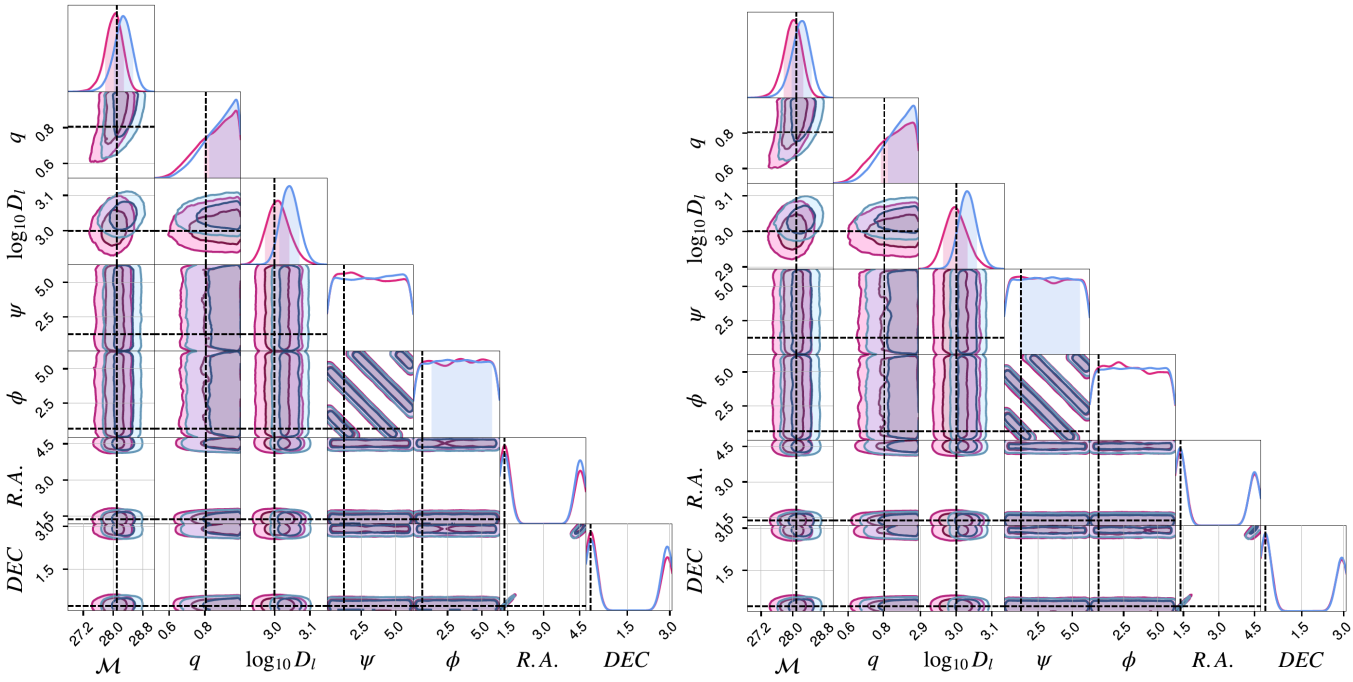
FIG. 9: Cornerplots using Gaussian likelihoods for **(a)** *known noise* (grey color) and **(b)** *fitting for the (unknown) noise* (light-blue color) for case in table I).

for the two different case studies of Sect. III C 1. The injections were intentionally chosen to overlap in time and frequency in order to stress-test the robustness of the likelihood under non-ideal noise conditions.

Fig. 10 shows the posterior distributions for the targeted BBH parameters obtained in the presence of multiple overlapping injections. The $\mathcal{H}$ and $\mathcal{W}$ likelihoods are shown for comparison. While the overall structure of the posteriors is almost similar, the $\mathcal{H}$ likelihood exhibits increased robustness against contamination from unresolved signals, whereas the $\mathcal{W}$ likelihood displays a mild tendency toward biased inference.

Fig. 11 and 12 demonstrate that the $\mathcal{H}$ likelihood provides a stable and accurate reconstruction, with the median reconstructed almost perfectly overlapping with the injected signal. The Gaussian likelihood (which is the one that most GW pipelines use) fails to recover the parameters of the injected signal, with significant portions of the injected waveform lying outside the inferred credible region. The $\mathcal{W}$ likelihood captures the injected signal within the 90% credible interval.

Despite the substantially different configurations of overlapping sub-threshold BBH signals, the resulting posterior distributions are strikingly similar in both experiments. This robustness indicates that the inference is primarily driven by the common high-SNR BBH signal and the common (real) noise, while the contribution of unresolved sub-threshold signals is effectively incorpo-



(a) Long-lasting BBH (see Table VI).

(b) Multiple injections of BBHs (see Table VII).

FIG. 10: Corner plot using the $\mathcal{W}$ (blue) and the $\mathcal{H}$ (pink) likelihoods. The PE refers to the targeted BBH (see Table III) when we simultaneously inject (a) a long-lasting BBH and (b) multiple low-SNR BBHs.

rated into the noise model.

Parameter	**Value**
$\mathrm{SNR}_{\mathrm{L1}}$	3.42
$\mathrm{SNR}_{\mathrm{H1}}$	3.39
Mass 1, m_1 [$M_\odot$]	15
Mass 2, m_2 [$M_\odot$]	8
Distance, D [Mpc]	2000
Right Ascension (RA) [rad]	2.75
Declination (Dec) [rad]	1.79
Polarization, ψ [rad]	0.60
Phase at coalescence, ϕ_0 [rad]	3.91
Inclination angle [rad]	0.40
Primary object tilt [rad]	0
Secondary object tilt [rad]	0
Primary dimensionless spin magnitude	0
Secondary dimensionless spin magnitude	0
Relative spin azimuthal angle, ϕ_{12} [rad]	0
Spin phase angle, ϕ_{jl} [rad]	0

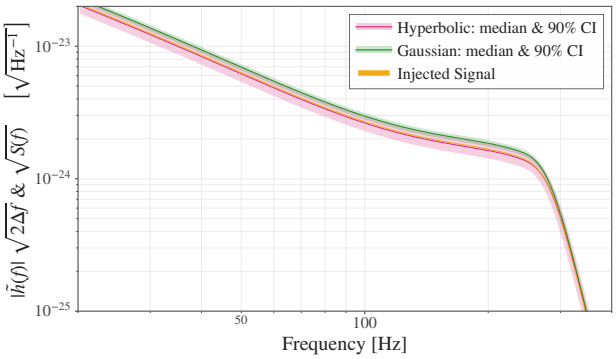
TABLE VI: Parameters of the long-lasting BBH injection in LIGO data.

These results suggest that the PSD correction inferred through the Whittle and Hyperbolic likelihoods effectively regularizes the inference problem by absorbing the contribution of unresolved BBH signals into an effective noise floor. As a consequence, the identifiability of the high-SNR BBH parameters is preserved even in the presence of complex signal overlap. Moreover, the improved stability of the $\mathcal{H}$ formulation highlights its enhanced ability to capture residual non-Gaussian features and PSD mismodeling effects, which become particularly relevant in scenarios dominated by overlapping subthreshold signals. In regimes characterized by stronger non-Gaussian noise features, the $\mathcal{W}$ likelihood tends to produce increasingly biased posterior estimates, indicating a breakdown of the Gaussian approximation underlying the Whittle formulation (as shown in Figures 6, 7, 8 and 13).

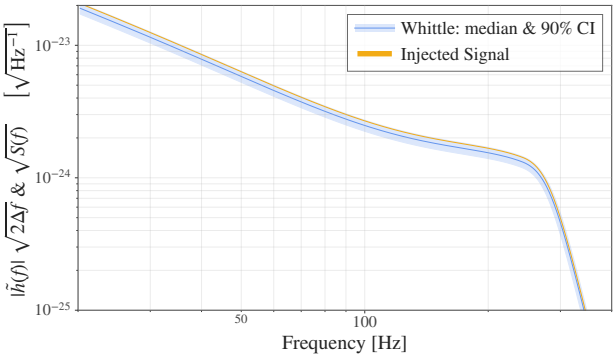
Fig. 13 shows the posterior distributions for the targeted BBH parameters obtained in the presence of the blip "glitch". In this glitch-dominated regime, the Whittle likelihood—which assumes Gaussian stationary noise—struggles to disentangle signal and noise contributions, leading to biased results, especially for the $\mathcal{M}$ and distance. Importantly, the Hyperbolic posteriors remain consistent with the injected parameter values across most dimensions. Overall (and given the reconstruction plots 6 and 7), the Hyperbolic likelihood provides a more reliable and interpretable posterior structure in the presence of glitches, making it particularly suitable for real detector data.

#	$\mathbf{SNR_{L1}}$	$\mathbf{SNR_{H1}}$	$\boldsymbol{\mathcal{M}}$ $[M_\odot]$	**q**	$\log_{\mathbf{10}}\mathbf{D_l}$ [Mpc]	$\boldsymbol{\psi}$ [rad]	$\boldsymbol{\phi_0}$ [rad]	**ra** [rad]	**dec** [rad]
1	4.4	2.6	17.37	0.91	3.46	1.93	1.99	4.43	-2.84
2	5.4	4.4	18.61	0.86	3.46	0.24	3.57	1.37	1.10
3	1.9	3.4	18.04	0.87	3.46	1.16	5.46	5.81	0.60
4	3.1	2.6	17.39	0.70	3.45	2.93	2.74	2.78	0.21
5	1.3	2.4	17.76	0.73	3.47	2.05	5.04	5.71	-2.87
6	4.8	4.2	17.26	0.86	3.47	1.25	0.90	0.38	0.39
7	2.6	2.4	17.34	0.87	3.46	2.48	4.42	1.16	-1.07

TABLE VII: Parameters of the low SNR BBH injections.

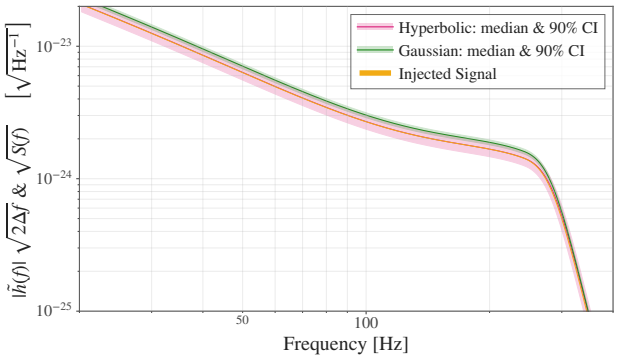


(a) Hyperbolic (pink) and Gaussian (green) likelihood

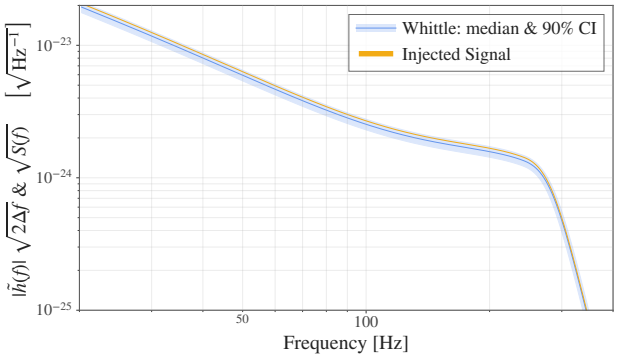


(b) Whittle likelihood

FIG. 11: Reconstruction (median and 90% C.I. reported) in the frequency domain of the targeted BBH merger in H1 detector. This case refers to the data with the long-lasting BBH injected. The reconstrucion presented concerns the PE by adopting the *(a)* $\mathcal{H}$ and Gaussian and *(b)* the $\mathcal{W}$ likelihood.



(a) Hyperbolic (pink) and Gaussian (green) likelihood



(b) Whittle likelihood

FIG. 12: Similar to Fig. 11, but for the case with the seven low-SNR BBH injections.

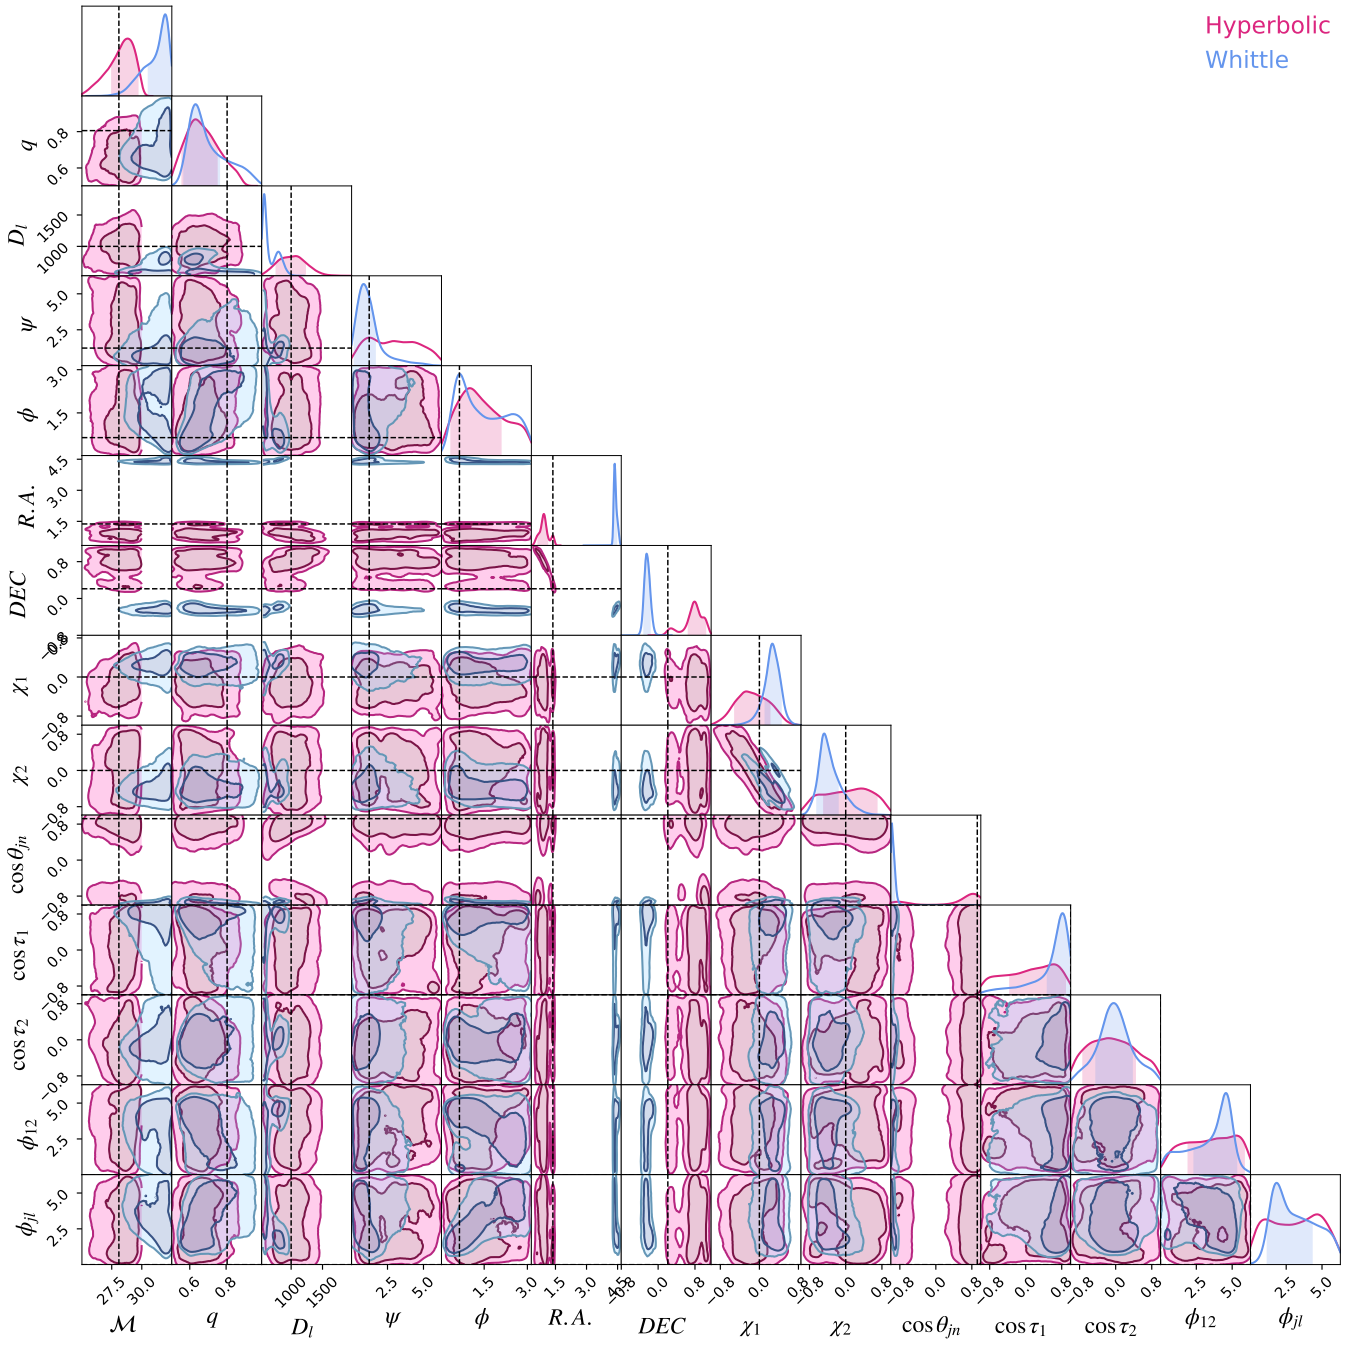


FIG. 13: Corner plot using the $\mathcal{W}$ (blue) and the $\mathcal{H}$ (pink) likelihood. This case corresponds to the case of the dataset with the "blip" glitch. The parameters of the injected BBH are given in Table III.

[1] T. B. Littenberg and N. J. Cornish, Prototype global analysis of lisa data with multiple source types, Phys. Rev. D **107**, 063004 (2023).

[2] N. Christensen and R. Meyer, Parameter estimation with gravitational waves, Rev. Mod. Phys. **94**, 025001 (2022).

[3] T. B. Littenberg and N. J. Cornish, Bayesian inference for spectral estimation of gravitational wave detector noise, Phys. Rev. D **91**, 084034 (2015), arXiv:1410.3852 [gr-qc].

[4] Q. Baghi, N. Karnesis, J.-B. Bayle, M. Besançon, and H. Inchauspé, Uncovering gravitational-wave backgrounds from noises of unknown shape with lisa, Journal of Cosmology and Astroparticle Physics **2023** (04), 066.

[5] L. Speri, N. Karnesis, A. I. Renzini, and J. R. Gair, A roadmap of gravitational wave data analysis, Nature Astron. **6**, 1356 (2022).

[6] A. Sasli, N. Karnesis, and N. Stergioulas, Heavy-tailed likelihoods for robustness against data outliers: Applications to the analysis of gravitational wave data, Phys. Rev. D **108**, 103005 (2023), arXiv:2305.04709 [gr-qc].

[7] A. Sasli, *Robust data analysis methods in gravitational wave astronomy*, Ph.D. thesis, Aristotle University of Thessaloniki (2025).

[8] M. Armano *et al.* (LISA Pathfinder), Transient acceleration events in LISA Pathfinder data: Properties and possible physical origin, Phys. Rev. D **106**, 062001 (2022), arXiv:2205.11938 [astro-ph.IM].

[9] Q. Baghi, N. Korsakova, J. Slutsky, E. Castelli, N. Karnesis, and J.-B. Bayle, Detection and characterization of instrumental transients in LISA Pathfinder and their projection to LISA, Phys. Rev. D **105**, 042002 (2022), arXiv:2112.07490 [gr-qc].

[10] R. Abbott, et al., Gwtc-3: Compact binary coalescences observed by ligo and virgo during the second part of the third observing run, Physical Review X **13**, 10.1103/physrevx.13.041039 (2023).

[11] Abbott R., et al., Gwtc-2.1: Deep extended catalog of compact binary coalescences observed by ligo and virgo during the first half of the third observing run (2021).

[12] R. Udall, S. Bini, K. Chatziioannou, D. Davis, S. Hourihane, Y. Lecoeuche, J. McIver, and S. Miller, Inferring the spins of merging black holes in the presence of data-quality issues (2025), arXiv:2510.05029 [gr-qc].

[13] K. Dey, N. Karnesis, A. Toubiana, E. Barausse, N. Korsakova, Q. Baghi, and S. Basak, Effect of data gaps on the detectability and parameter estimation of massive black hole binaries with LISA, Phys. Rev. D **104**, 044035 (2021), arXiv:2104.12646 [gr-qc].

[14] Q. Baghi, I. Thorpe, J. Slutsky, J. Baker, T. Dal Canton, N. Korsakova, and N. Karnesis, Gravitational-wave parameter estimation with gaps in LISA: a Bayesian data augmentation method, Phys. Rev. D **100**, 022003 (2019), arXiv:1907.04747 [gr-qc].

[15] S. Drasco and E. E. Flanagan, Detection methods for non-gaussian gravitational wave stochastic backgrounds, Phys. Rev. D **67**, 082003 (2003).

[16] S. Wu and A. H. Nitz, Mock data study for next-generation ground-based detectors: The performance loss of matched filtering due to correlated confusion noise, Physical Review D **107**, 10.1103/physrevd.107.063022 (2023).

[17] A. D. Johnson, K. Chatziioannou, and W. M. Farr, Source confusion from neutron star binaries in ground-based gravitational wave detectors is minimal, Physical Review D **109**, 10.1103/physrevd.109.084015 (2024).

[18] Abbott R., et al., Open data from the third observing run of ligo, virgo, kagra, and geo, The Astrophysical Journal Supplement Series **267**, 29 (2023).

[19] S. Ghonge, J. Brandt, J. M. Sullivan, M. Millhouse, K. Chatziioannou, J. A. Clark, T. Littenberg, N. Cornish, S. Hourihane, and L. Cadonati, Assessing and mitigating the impact of glitches on gravitational-wave parameter estimation: A model agnostic approach, Phys. Rev. D **110**, 122002 (2024).

[20] R. Macas, J. Pooley, L. K. Nuttall, D. Davis, M. J. Dyer, Y. Lecoeuche, J. D. Lyman, J. McIver, and K. Rink, Impact of noise transients on low latency gravitational-wave event localization, Phys. Rev. D **105**, 103021 (2022).

[21] B. Zackay, T. Venumadhav, J. Roulet, L. Dai, and M. Zaldarriaga, Detecting gravitational waves in data with non-stationary and non-gaussian noise, Physical Review D **104**, 10.1103/physrevd.104.063034 (2021).

[22] M. C. Edwards, P. Maturana-Russel, R. Meyer, J. Gair, N. Korsakova, and N. Christensen, Identifying and addressing nonstationary lisa noise, Phys. Rev. D **102**, 084062 (2020).

[23] Aasi J., et al., Advanced ligo, Classical and Quantum Gravity **32**, 074001 (2015).

[24] Acernese F., et al., Advanced virgo: a second-generation interferometric gravitational wave detector, Classical and Quantum Gravity **32**, 024001 (2014).

[25] Y. Aso, Y. Michimura, K. Somiya, M. Ando, O. Miyakawa, T. Sekiguchi, D. Tatsumi, and H. Yamamoto, Interferometer design of the KAGRA gravitational wave detector, Phys. Rev. D **88**, 043007 (2013), arXiv:1306.6747 [gr-qc].

[26] Punturo M., et al, The einstein telescope: a third-generation gravitational wave observatory, Classical and Quantum Gravity **27**, 194002 (2010).

[27] V. Srivastava, D. Davis, K. Kuns, P. Landry, S. Ballmer, M. Evans, E. D. Hall, J. Read, and B. S. Sathyaprakash, Science-driven tunable design of cosmic explorer detectors, The Astrophysical Journal **931**, 22 (2022).

[28] Ackley K., et al.,, Neutron Star Extreme Matter Observatory: A kilohertz-band gravitational-wave detector in the global network, Publications of the Astronomical Society of Australia **37**, e047 (2020), arXiv:2007.03128 [astro-ph.HE].

[29] Amaro-Seoane P., et al., Laser interferometer space antenna (2017).

[30] Y. Gong, J. Luo, and B. Wang, Concepts and status of Chinese space gravitational wave detection projects, Nature Astronomy **5**, 881 (2021), arXiv:2109.07442 [astro-ph.IM].

[31] H. Wang, D. MingHui, P. Xu, and Z. Yu-Feng, Challenges in space-based gravitational wave data analysis and applications of artificial intelligence, Scientia Sinica Physica Mechanica & Astronomica **54**, 270403 (2024).

[32] B. P. Abbott, others, T. L. S. Collaboration, and the Virgo Collaboration, A guide to ligo-virgo detector noise and extraction of transient gravitational-wave signals, Classical and Quantum Gravity **37**, 055002 (2020).

[33] N. J. Cornish and T. B. Littenberg, BayesWave:

Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches, Class. Quant. Grav. **32**, 135012 (2015), arXiv:1410.3835 [gr-qc].
[34] R. Umstätter, N. Christensen, M. Hendry, R. Meyer, V. Simha, J. Veitch, S. Vigeland, and G. Woan, Bayesian modeling of source confusion in lisa data, Phys. Rev. D **72**, 022001 (2005).
[35] T. B. Littenberg, N. J. Cornish, K. Lackeos, and T. Robson, Global analysis of the gravitational wave signal from galactic binaries, Phys. Rev. D **101**, 123021 (2020).
[36] T. B. Littenberg and N. J. Cornish, Prototype global analysis of lisa data with multiple source types, Phys. Rev. D **107**, 063004 (2023).
[37] Q. Baghi, N. Karnesis, J.-B. Bayle, M. Besançon, and H. Inchauspé, Uncovering gravitational-wave backgrounds from noises of unknown shape with lisa, Journal of Cosmology and Astroparticle Physics **2023** (04), 066.
[38] M. C. Edwards, P. Maturana-Russel, R. Meyer, J. Gair, N. Korsakova, and N. Christensen, Identifying and Addressing Nonstationary LISA Noise, Phys. Rev. D **102**, 084062 (2020), arXiv:2004.07515 [gr-qc].
[39] M. C. Edwards, R. Meyer, and N. Christensen, Bayesian semiparametric power spectral density estimation with applications in gravitational wave data analysis, Phys. Rev. D **92**, 064011 (2015).
[40] L. Martellini and T. Regimbau, Semiparametric approach to the detection of non-Gaussian gravitational wave stochastic backgrounds, Phys. Rev. D **89**, 124009 (2014), arXiv:1405.5775 [astro-ph.CO].
[41] S. Hamimeche and A. Lewis, Likelihood Analysis of CMB Temperature and Polarization Power Spectra, Phys. Rev. D **77**, 103013 (2008), arXiv:0801.0554 [astro-ph].
[42] L. Verde, H. V. Peiris, D. N. Spergel, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright, First-year wilkinson microwave anisotropy probe (wmap)* observations: Parameter estimation methodology, The Astrophysical Journal Supplement Series **148**, 195 (2003).
[43] R. Flauger, N. Karnesis, G. Nardini, M. Pieroni, A. Ricciardone, and J. Torrado, Improved reconstruction of a stochastic gravitational wave background with lisa, Journal of Cosmology and Astroparticle Physics **2021** (01), 059–059.
[44] J. D. E. Creighton, Data analysis strategies for the detection of gravitational waves in non-gaussian noise, Phys. Rev. D **60**, 021101 (1999).
[45] B. Allen, J. D. E. Creighton, E. E. Flanagan, and J. D. Romano, Robust statistics for deterministic and stochastic gravitational waves in non-gaussian noise: Frequentist analyses, Phys. Rev. D **65**, 122002 (2002).
[46] R. Legin, A. Adam, Y. Hezaveh, and L. Perreault-Levasseur, Beyond gaussian noise: A generalized approach to likelihood analysis with non-gaussian noise, The Astrophysical Journal Letters **949**, L41 (2023).
[47] N. Karnesis, A. Sasli, R. Buscicchio, and N. Stergioulas, Characterization of non-gaussian stochastic signals with heavier-tailed likelihoods (2024), arXiv:2410.14354 [gr-qc].
[48] E. Eberlein and K. Prause, The generalized hyperbolic model: financial derivatives and risk measures, in *Mathematical Finance—Bachelier Congress 2000* (Springer, 2002) pp. 245–267.
[49] K. Prause, The generalized hyperbolic model: Estimation, financial derivatives, and risk measures, PhD thesis (1999).
[50] Q. Baghi, N. Karnesis, J.-B. Bayle, M. Besançon, and H. Inchauspé, Uncovering gravitational-wave backgrounds from noises of unknown shape with lisa, Journal of Cosmology and Astroparticle Physics **2023** (04), 066.
[51] M. Mapelli, Astrophysics of stellar black holes (2018), arXiv:1809.09130 [astro-ph.HE].
[52] Abbott B.P., et al., Observation of Gravitational Waves from a Binary Black Hole Merger, Phys. Rev. Lett. **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
[53] Abbott B. P., et al. (LIGO Scientific Collaboration and Virgo Collaboration), Binary black hole mergers in the first advanced ligo observing run, Phys. Rev. X **6**, 041015 (2016).
[54] Abbott B. P., et al. (LIGO Scientific Collaboration and Virgo Collaboration), Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs, Phys. Rev. X **9**, 031040 (2019).
[55] Abbott R., et al. (LIGO Scientific Collaboration and Virgo Collaboration), Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run, Phys. Rev. X **11**, 021053 (2021).
[56] Abbott R., et al., The population of merging compact binaries inferred using gravitational waves through gwtc-3 (2021).
[57] A. H. Nitz, S. Kumar, Y.-F. Wang, S. Kastha, S. Wu, M. Schäfer, R. Dhurkunde, and C. D. Capano, 4-ogc: Catalog of gravitational waves from compact binary mergers, The Astrophysical Journal **946**, 59 (2023).
[58] A. E. Koloniari, E. C. Koursoumpa, P. Nousi, P. Lampropoulos, N. Passalis, A. Tefas, and N. Stergioulas, New gravitational wave discoveries enabled by machine learning (2024), arXiv:2407.07820 [gr-qc].
[59] Abbott, B. P., et al., A gravitational-wave measurement of the hubble constant following the second observing run of advanced ligo and virgo, The Astrophysical Journal **909**, 218 (2021).
[60] Abbott, B. P., et al. (LIGO Scientific Collaboration and Virgo Collaboration), Binary black hole mergers in the first advanced ligo observing run, Phys. Rev. X **6**, 041015 (2016).
[61] Abbott, R., et al. (LIGO Scientific Collaboration and Virgo Collaboration), Gw190412: Observation of a binary-black-hole coalescence with asymmetric masses, Phys. Rev. D **102**, 043015 (2020).
[62] Abbott, B. P., et al., Binary black hole population properties inferred from the first and second observing runs of advanced ligo and advanced virgo, The Astrophysical Journal Letters **882**, L24 (2019).
[63] Abbott, B. P., et al., The rate of binary black hole mergers inferred from advanced ligo observations surrounding gw150914, The Astrophysical Journal Letters **833**, L1 (2016).
[64] V. Kalogera, C. P. L. Berry, M. Colpi, S. Fairhurst, S. Justham, I. Mandel, A. Mangiagli, M. Mapelli, C. Mills, B. S. Sathyaprakash, R. Schneider, T. Tauris, and R. Valiante, Deeper, wider, sharper: Next-generation ground-based gravitational-wave observations of binary black holes (2019), arXiv:1903.09220 [astro-ph.HE].
[65] W. M. Farr, M. Fishbach, J. Ye, and D. E. Holz, A future percent-level measurement of the hubble expansion at redshift 0.8 with advanced ligo, The Astrophysical Jour-

nal Letters **883**, L42 (2019).
[66] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis*, 2nd ed. (Chapman and Hall/CRC, 2004).
[67] L. Barack and C. Cutler, Lisa capture sources: Approximate waveforms, signal-to-noise ratios, and parameter estimation accuracy, Phys. Rev. D **69**, 082005 (2004).
[68] C. Cutler and E. E. Flanagan, Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform?, Phys. Rev. D **49**, 2658 (1994).
[69] N. J. Cornish, T. B. Littenberg, B. Bécsy, K. Chatziioannou, J. A. Clark, S. Ghonge, and M. Millhouse, BayesWave analysis pipeline in the era of gravitational wave observations, Phys. Rev. D **103**, 044006 (2021), arXiv:2011.09494 [gr-qc].
[70] P. Auclair *et al.* (LISA Cosmology Working Group), Cosmology with the Laser Interferometer Space Antenna, Living Rev. Rel. **26**, 5 (2023), arXiv:2204.05434 [astro-ph.CO].
[71] C. R. Contaldi, M. Pieroni, A. I. Renzini, G. Cusin, N. Karnesis, M. Peloso, A. Ricciardone, and G. Tasinato, Maximum likelihood map-making with the Laser Interferometer Space Antenna, Phys. Rev. D **102**, 043502 (2020), arXiv:2006.03313 [astro-ph.CO].
[72] M. Armano *et al.*, Calibrating the system dynamics of LISA Pathfinder, Phys. Rev. D **97**, 122002 (2018), arXiv:1806.08581 [astro-ph.IM].
[73] C. Röver, Student-t based filter for robust signal detection, Phys. Rev. D **84**, 122004 (2011).
[74] C. Röver, R. Meyer, and N. Christensen, Modelling coloured residual noise in gravitational-wave signal processing, Classical and Quantum Gravity **28**, 015010 (2011).
[75] M. Tinto and S. V. Dhurandhar, TIME DELAY, Living Rev. Rel. **8**, 4 (2005), arXiv:gr-qc/0409034.
[76] S. Babak, A. Petiteau, and M. Hewitson, LISA Sensitivity and SNR Calculations (2021), arXiv preprint, arXiv:2108.01167 [astro-ph.IM].
[77] N. Karnesis, M. L. Katz, N. Korsakova, J. R. Gair, and N. Stergioulas, Eryn: a multipurpose sampler for bayesian inference, Monthly Notices of the Royal Astronomical Society **526**, 4814–4830 (2023).
[78] M. Karamanis, D. Nabergoj, F. Beutler, J. A. Peacock, and U. Seljak, pocomc: A python package for accelerated bayesian inference in astronomy and cosmology (2022), arXiv:2207.05660 [astro-ph.IM].
[79] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, `emcee`: The MCMC hammer, Publications of the Astronomical Society of the Pacific **125**, 306 (2013).
[80] M. Karamanis and U. Seljak, Persistent Sampling: Enhancing the Efficiency of Sequential Monte Carlo, Statistics and Computing **35**, 1 (2025), arXiv:2407.20722.
[81] M. J. Williams, M. Karamanis, Y. Luo, and U. Seljak, Validating sequential monte carlo for gravitational-wave inference, Monthly Notices of the Royal Astronomical Society **543**, 1479–1493 (2025).
[82] S. Vretinaris, G. Vretinaris, C. Mermigkas, M. Karamanis, and N. Stergioulas, Robust and fast parameter estimation for gravitational waves from binary neutron star merger remnants, Physical Review D **113**, 10.1103/g1qs-j74x (2026).
[83] The LISA Data Challenges, `https://lisa-ldc.lal.in2p3.fr/` (2023), accessed: 2023-01-01.
[84] M. L. Katz, S. Marsat, A. J. Chua, S. Babak, and S. L. Larson, GPU-accelerated massive black hole binary parameter estimation with LISA, Physical Review D **102**, 10.1103/physrevd.102.023033 (2020).
[85] M. L. Katz, Fully automated end-to-end pipeline for massive black hole binary signal extraction from LISA data, Physical Review D **105**, 10.1103/physrevd.105.044055 (2022).
[86] M. Katz, mikekatz04/bbhx: First official public release (2021).
[87] S. Khan, S. Husa, M. Hannam, L. London, M. Pürrer, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. Phenomenological model for the 22-mode, Phys. Rev. D **93**, 044007 (2016), arXiv:1508.07253 [gr-qc].
[88] M. Tinto and J. W. Armstrong, Cancellation of laser noise in an unequal-arm interferometer detector of gravitational radiation, Phys. Rev. D **59**, 102003 (1999).
[89] T. A. Prince, M. Tinto, S. L. Larson, and J. W. Armstrong, Lisa optimal sensitivity, Phys. Rev. D **66**, 122002 (2002).
[90] S. Babak, M. Hewitson, and A. Petiteau, Lisa sensitivity and snr calculations (2021).
[91] R. Flauger, N. Karnesis, G. Nardini, M. Pieroni, A. Ricciardone, and J. Torrado, Improved reconstruction of a stochastic gravitational wave background with lisa, Journal of Cosmology and Astroparticle Physics **2021** (01), 059.
[92] NVIDIA, P. Vingelmann, and F. H. Fitzek, Cuda, release: 10.2.89 (2020).
[93] G. Ashton *et al.*, BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy, Astrophys. J. Suppl. **241**, 27 (2019), arXiv:1811.02042 [astro-ph.IM].
[94] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. J. Forteza, and A. Bohé, Frequency-domain gravitational waves from nonprecessing black-hole binaries. ii. a phenomenological model for the advanced detector era, Phys. Rev. D **93**, 044007 (2016).
[95] J. Robnik and U. Seljak, Matched filtering with non-gaussian noise for planet transit detections, Monthly Notices of the Royal Astronomical Society **504**, 5829 (2021), https://academic.oup.com/mnras/article-pdf/504/4/5829/38103994/stab1178.pdf.
[96] J. Robnik, U. Seljak, J. M. Jenkins, and S. Bryson, Exoplanet transit search at the detection limit: detection and false alarm vetting pipeline (2026), arXiv:2601.07465 [astro-ph.EP].
[97] Https://github.com/asasli/HyperWave/.
[98] Https://github.com/asasli/HyperWave/.
[99] C. D. Manning and H. Schütze, *Foundations of Statistical Natural Language Processing* (MIT Press, Cambridge, MA, USA, 1999).