DRAFT VERSION MAY 22, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

GWTC-5.0: Methods for Identifying and Characterizing Gravitational-wave Transients

THE LIGO SCIENTIFIC COLLABORATION, THE VIRGO COLLABORATION, AND THE KAGRA COLLABORATION

ABSTRACT

The Gravitational-Wave Transient Catalog (GWTC) is a collection of candidate gravitational-wave transient signals identified and characterized by the LIGO–Virgo–KAGRA Collaboration. Producing the contents of the GWTC from detector data requires complex analysis methods. These comprise techniques to model the signal; identify the transients in the data; evaluate the quality of the data and mitigate possible instrumental issues; infer the parameters of each transient; compare the data with the waveform models for compact binary coalescences, and handle the large amount of results associated with all these different analyses. In this paper, we describe the methods employed to produce the catalog's fifth release, GWTC-5.0, focusing on the analysis of the second part of the fourth observing run of LIGO, Virgo and KAGRA.

Keywords: Gravitational wave astronomy (675); Gravitational wave detectors (676); Gravitational wave sources (677); Stellar mass black holes (1611); Neutron stars (1108)

1. INTRODUCTION

Interferometric gravitational-wave (GW) detectors produce a calibrated discrete digital time series $d(t)$ known as the *strain* (a dimensionless measure of the relative difference in arm length of the interferometers). The Laser Interferometer Gravitational-Wave Observatory (LIGO; Aasi et al. 2015) and Virgo (Acernese et al. 2015) detectors are the most sensitive to date. Alongside the developing KAGRA detector (Akutsu et al. 2019) the LIGO–Virgo–KAGRA Collaboration (LVK) has recently undertaken the fourth observing run (O4). However, the data produced are dominated by detector noise, with only occasional occurrences of detectable transient GW signals (Abbott et al. 2020; Abbott et al. 2020a); to date, all such observed signals probably arise from compact binary coalescences (CBCs) involving black holes (BHs) and neutron stars (NSs). This paper describes the methodology used to analyze the calibrated strain data up to 2025 January 28, the end date of the second part of the fourth observing run (O4b), and produce version 5.0 of the Gravitational Wave Transient Catalog (GWTC), hereafter referred to as GWTC-5.0. This catalog release also includes updated search results and parameter estimation of additional events from first part of the fourth observing run (O4a). To clarify that these O4a results supersede those in GWTC-4.0, we name them GWTC-4.1. However, as GWTC-5.0 is a cumulative catalog, all the results labeled GWTC-4.1 are included in GWTC-5.0 by definition. For a general introduction to GWTC-5.0, see Abac et al. (2026a) which also contains a description of the observed source classes and data-analysis nomenclature which should be read as a background to this methodology paper. The scientific results of GWTC-5.0 are presented in Abac et al. (2026b). Data-analysis methods not directly related to producing results for GWTC-5.0, such as searches for continuous GWs or subsolar-mass CBCs, will be described elsewhere.

There are many interconnected elements to the data processing methodology described in this work. To provide a visual guide and summary, Figure 1 shows a diagram of the data-processing workflow. In Section 2, we introduce the fundamental concepts behind modeling GW waveforms from CBC sources and describe the waveform approximants used in later analyses. The data analysis process then starts with the calibrated strain data $d(t)$ and associated auxiliary data, which is produced by the LVK detectors and is the input to the analysis. The strain data and auxiliary data are inputs to the process of searching for signals and compiling a list of candidates from the strain data which we describe in Section 3. Section 4 discusses how we assess data quality around candidates and mitigate the impact of potential instrumental issues. The strain data, candidate lists, and data-quality information are then inputs into the methods used to infer the properties of the signals and their sources which we describe in Section 5. In Section 6 we detail consistency tests performed on selected candidates to evaluate how well CBC waveform models match the data. Section 7 describes the technologies used to manage the flow of informa-

Corresponding author: LSC P&P Committee, via LVK Publications as proxy
lvc.publications@ligo.org

tion throughout the analysis process as shown in Figure 1. Finally, we conclude in Section 8.

2. MODELING COMPACT BINARY COALESCENCE SIGNALS

In this section, we will describe the modeling of CBC waveform signals, including binary BH (BBH), binary NS (BNS), and NS–BH binary (NSBH) systems. This modeling is a crucial aspect for the detection and astrophysical interpretation of these systems and is used by modeled search algorithms (described in Section 3) and parameter estimation (PE; described in Section 5). The focus of this section is on the waveform models used in GWTC-5.0, although we include models employed in analyses of previous GWTC versions that have been superseded in newer versions, in particular the PE analyses performed for GWTC-1.0 and GWTC-2.0. For an introduction to waveform modeling in general and definitions of our coordinate conventions, see Section 5 of Abac et al. (2026a).

Several different modeling approaches have been followed in the development of CBC waveform models for the complete inspiral–merger–ringdown (IMR) stages of the signal. Inspiral-only models have been developed based on post-Newtonian (PN) theory (Blanchet 2014), using different PN-expansions of the energy balance equations for the two-body dynamics, under the TAYLOR family (Buonanno et al. 2009) for non-spinning systems and the SPINTAYLOR family (Sturani 2015; Isoyama et al. 2020) for spinning systems. The IMRPHENOM approach (Ajith et al. 2007) focuses on the description of the GW signal, traditionally in Fourier domain for an efficient implementation in the data analysis pipelines, combining PN and numerical relativity (NR) information into closed-form expressions for describing the inspiral, merger and ringdown stages of the signal. The effective-one-body (EOB) approach (Buonanno & Damour 1999, 2000) focuses on accurately describing the dynamics and resulting waveform of the system in the time-domain, using a combination of resummed analytical information and calibration to NR. Inside this approach, two main different development efforts have been followed, the SEOBNR (Buonanno et al. 2007) and TEOB (Damour & Nagar 2014; Nagar & Shah 2016) approaches. The NRSURROGATE approach (Blackman et al. 2017a) focuses on producing efficient surrogate models of the NR data, with or without hybridization with analytical waveforms, to deliver highly accurate waveforms in a validity region limited by the input numerical data, both in the number of cycles and in the coverage of parameter space. These modeling approaches provide the waveform models described in this section.

The majority of waveform models that we describe in this section correspond to CBC systems on quasi-circular orbits (also called quasi-spherical orbits when spin precession is present), therefore neglecting orbital eccentricity effects, since at the moment all models employed in the analysis of GWTC candidates have this restriction. One of the main reasons for this limitation is that mature eccentric waveform models have only been developed recently, and reviewed versions were not yet ready at the time of starting the analyses described in this work. This limitation can lead to biases in the mass parameters (Martel & Poisson 1999; Lower et al. 2018; Lenon et al. 2020; O'Shea & Kumar 2023; Favata et al. 2022), and also influence the inferred spins (Lenon et al. 2020; Romero-Shaw et al. 2020a; O'Shea & Kumar 2023; Favata et al. 2022; Morras et al. 2025a,b). For high-mass systems, where few orbits are observed, neglecting eccentricity can also lead to incorrectly identifying spin-precessing effects due to possible degeneracies of both effects in this regime (Ramos-Buades et al. 2020; Calderón Bustillo et al. 2021; Romero-Shaw et al. 2023). During a binary inspiral, the energy loss from GW emission rapidly circularizes an eccentric orbit, generally reducing expectation of significant eccentricity (Peters 1964) by the time the signal enters the sensitive frequency band of the interferometers (Tucker & Will 2021). Rate estimations (Wen 2003; Samsing & Ramirez-Ruiz 2017; Gupte et al. 2024) constrain the fraction of events with observable eccentricity to a small percentage.

Nevertheless, recent independent analyses have shown evidence of nonzero eccentricity in a few previously detected signals (Romero-Shaw et al. 2020a; Gamba et al. 2023a; Gayathri et al. 2022; Gupte et al. 2024; Planas et al. 2025b; Morras et al. 2025b; Planas et al. 2025a; Jan et al. 2025; Romero-Shaw et al. 2025; Kacanja et al. 2025).

Therefore, we cannot exclude a priori that a few candidates in GWTC-5.0 may present nonzero eccentricity.

In the following subsections, we will describe chronologically the relevant waveform modeling efforts that lead to models employed in GWTC analyses. For consulting the specific set of waveform models employed in each GWTC version, we refer the reader to Table 1, while a summary of the physics of each model is displayed in Table 2. Their usage on specific pipelines will be described in the corresponding sections of this work.

2.1. *BBH Models*

Within the IMRPHENOM and SEOBNR families, the first complete IMR models calibrated to NR were produced for the dominant spin-weighted spherical harmonic multipole of nonspinning systems (Ajith et al. 2007; Buonanno et al. 2007), and then extended to aligned-spin systems (Ajith et al. 2011; Santamaría et al. 2010; Taracchini et al. 2014b). Improved versions were developed for GWTC analyses: IMRPHENOMD (Husa et al. 2016; Khan et al. 2016) and SEOBNRV4 (Bohé et al. 2017), increasing the amount of analytical information included in the models, the number and coverage of the calibration dataset, specific details of the model expressions and better accuracy with NR. A highly optimized version of SEOBNRV4 was produced to reduce its computational cost in the time domain, SEOBNRV4_OPT (Devine et al. 2016), and reduced-order model (ROM) techniques were applied to obtain a fast Fourier-domain version of the original model, SEOBNRV4_ROM (Bohé et al. 2017). Within the TEOB approach, a version of the model for the dominant multipole of BBH systems was developed, TEOBRESUMS (Nagar et al. 2018), including a post-adiabatic

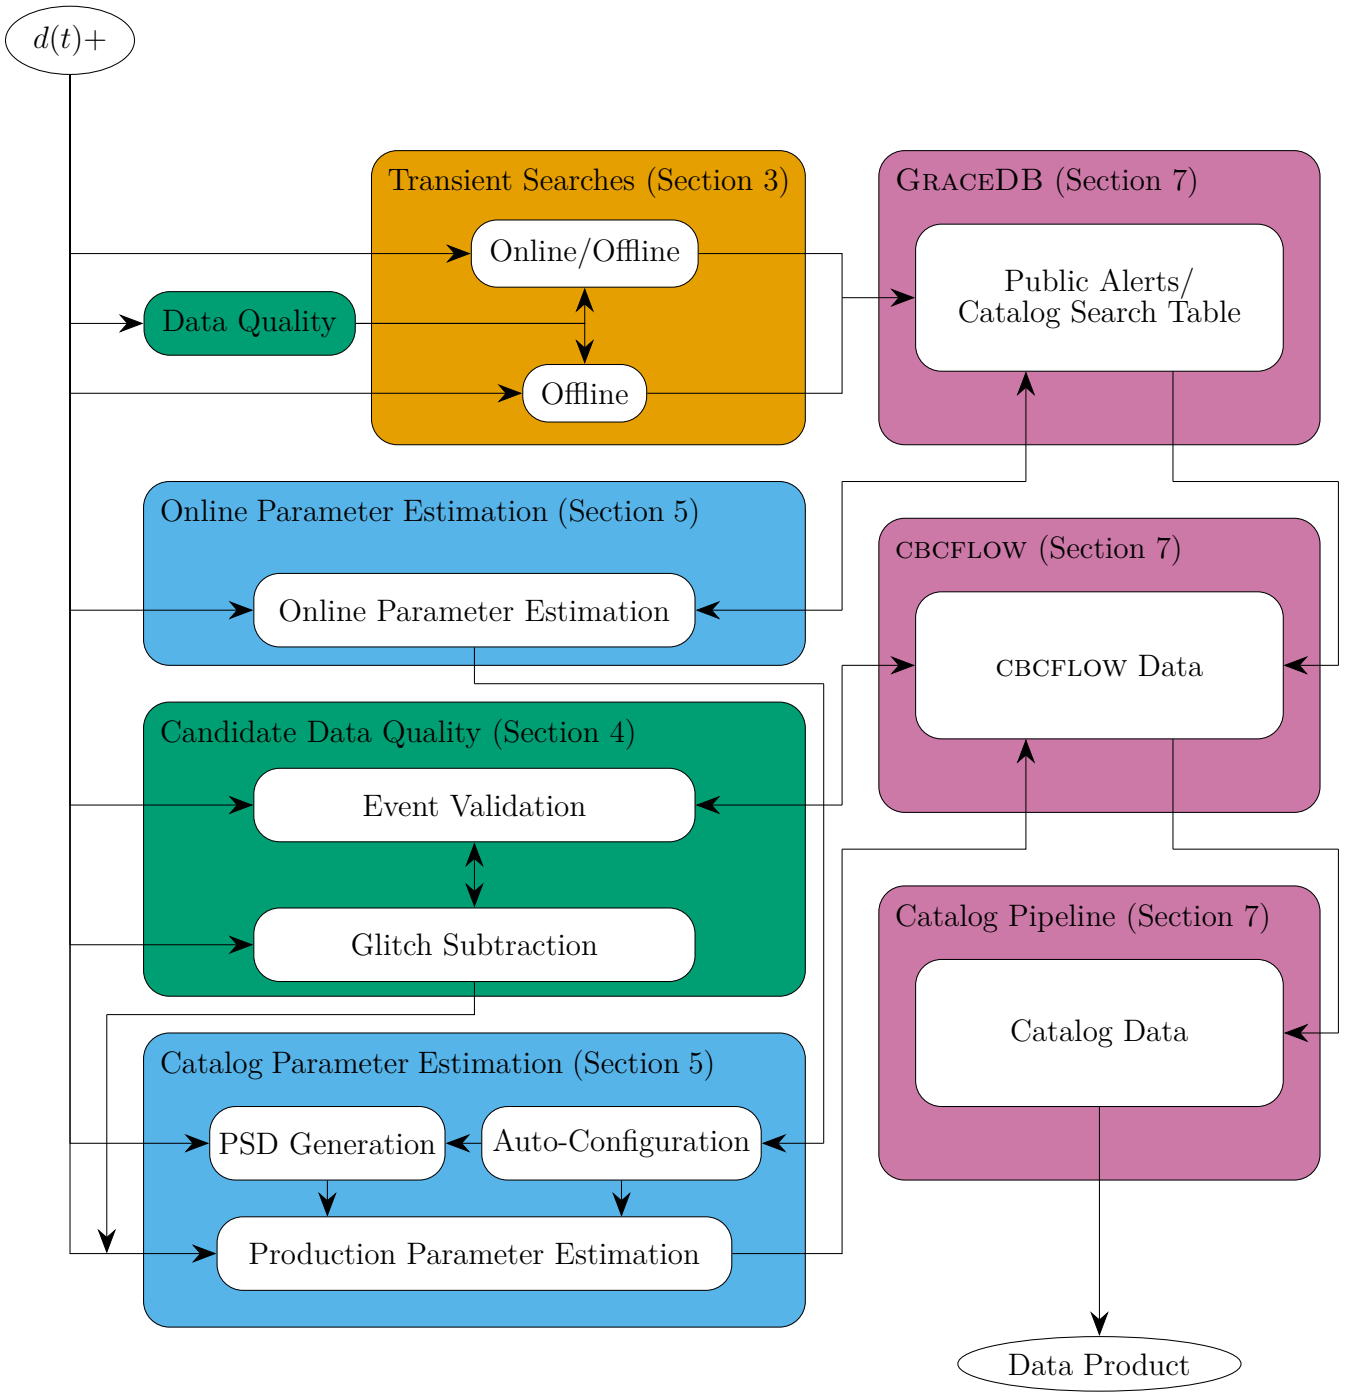

Figure 1. A high-level rendering of the flow of data from the strain and auxiliary data, denoted $d(t)+$, to the production of GWTC-5.0 described in this work. In this work, we use the terms *downstream* and *upstream* to refer to the flow of information in the analysis process (e.g. the parameter estimation is downstream of the compilation of search results since it depends on their outputs). This complex process enables the use of the most complete set of information for the final results while also leveraging preliminary studies to parallelise and reduce the overall analysis time. The term *online* (also referred to as *low-latency* in the literature) refers to analyses run on live data with a goal to upload candidates to GRACEDB immediately and enable rapid public alerts, while *offline* refers to analyses run with the goal to identify candidates for the GWTC. PSD refers to the power spectral density. In each box, we provide references to the sections of this paper where the methods are explained in more detail.

approximation for the dynamics that increases its computational efficiency (Nagar & Rettegno 2019). Additionally, the inspiral-only models TaylorF2 (Damour et al. 2001; Buonanno et al. 2009; Vines et al. 2011), which is an analytical Fourier-domain model, and SpinTaylorT4 (Sturani 2015; Isoyama et al. 2020), a time-domain model, have also been employed for searching GW signals.

Modeling spin-precession effects is crucial for precise measurements of spins (Vitale et al. 2014; Pratten et al. 2020b; Johnson-McDaniel et al. 2022b; Biscoveanu et al. 2021; Steinle & Kesden 2022), providing key information about the formation channels of the observed systems (Rodriguez et al. 2016; Stevenson et al. 2017; Talbot & Thrane 2017; Zhu et al. 2018), and breaking degeneracies in parameter inference (Vecchio 2004; Lang & Hughes 2006; Chatziioannou et al. 2015; Krishnendu & Ohme 2022). Spin-precession can significantly increase the complexity of the signal, and modeling efforts focused first on applying the *twisting-up* approach (Buonanno et al. 2003; Schmidt et al. 2011; Boyle et al. 2011; O'Shaughnessy et al. 2012; Schmidt et al. 2012) to the dominant-multipole models, differing on the description of the spin dynamics of the system. The IMRPhenom approach incorporated a closed-form solution of the next-to-next-to-leading order spin-precessing evolution equations for single-spin configurations (Marsat et al. 2013; Bohe et al. 2013), mapping then the expressions effectively to double-spin systems (Schmidt et al. 2015), and employing the twisting-up technique first to IMRPhenomC (Santamaría et al. 2010) to obtain IMRPhenomP (Hannam et al. 2014), and then to the more accurate model IMRPhenomD to obtain IMRPhenomPv2 (Hannam et al. 2014; Bohé et al. 2016). On the SEOBNR approach, spin-dynamics were incorporated via evolution of the EOB equations of motion for a quasi-circular spin-precessing Hamiltonian, constructing the spin-precessing polarizations from the quadrupolar multipoles in the co-precessing frame for producing SEOBNRv3 (Pan et al. 2014).

Subdominant harmonics in the signal were shown to be important for reducing several degeneracies in the analysis of signals (Capano et al. 2014; Graff et al. 2015; Abbott et al. 2021a, 2024), in particular in inclination–distance degeneracy and modeling efforts focused on their inclusion into current models. IMRPhenomHM (London et al. 2018) incorporated a set of the most important subdominant harmonics to the IMRPhenomD model, although without additional calibration of these harmonics to numerical data. SEOBNRv4HM (Cotesta et al. 2018) incorporated a similar list of harmonics to SEOBNRv4, with explicit calibration of the waveform modes to NR and test-particle waveforms, and provided an efficient ROM version in Fourier-domain, SEOBNRv4HM_ROM (Cotesta et al. 2020). Spin-precessing versions of these multipolar waveform models were developed in parallel, with some improvements in the spin-dynamics description in the IMRPhenom approach, resulting in IMRPhenomPv3HM (Khan et al. 2020) and SEOBNRv4PHM (Ossokine et al. 2020). Additionally, the NRSurrogate approach started producing the first surrogate models for multipolar spin-precessing signals, first with NRSur7dq2 (Blackman et al. 2017b) and then extending the parameter space coverage with NRSur7dq4 (Varma et al. 2019).

Inside the IMRPhenom approach, a new generation of waveform models was developed improving substantially the accuracy of the previous generation, providing a new model for the dominant harmonic of aligned-spin signal, IMRPhenomXAS (Pratten et al. 2020a), a model for subdominant harmonics explicitly calibrated to NR simulations, IMRPhenomXHM (García-Quirós et al. 2020) and its extension to spin-precessing signals through the twisting-up technique and the employment of the multiscale expression for the spin dynamics (Klein et al. 2013; Chatziioannou et al. 2013; Klein et al. 2014), IMRPhenomXPHM (Pratten et al. 2021). This generation also reduced substantially the computational cost of waveform generation via an implementation of the multibanding method (Vinciguerra et al. 2017; García-Quirós et al. 2021). On a parallel effort, a new phenomenological family of models was developed in the time domain, IMRPhenomTPHM (Estellés et al. 2021, 2022b,a), with the aim of overcoming the limitations of the stationary-phase approximation (SPA) in the modeling of precessing signals, and including for the first time in a phenomenological model accurate numerical solutions of the spin-precession equations.

Further improvements for spin-precessing signals were recently introduced in the IMRPhenomXPHM model, introducing for the first time explicit calibration with spin-precessing NR simulations in IMRPhenomXO4a (Hamilton et al. 2021; Thompson et al. 2024) as well as the inclusion of the dominant multipole equatorial asymmetry (Ghosh et al. 2024), a key effect for accurately describing systems with large remnant recoil (Varma et al. 2020; Borchers et al. 2024) and improving accuracy of spin-precessing models (Ramos-Buades et al. 2020a). This effect has been shown in several recent studies to be relevant for the correct inference of spin-precessing signals (Kolitsidou et al. 2024; Estellés et al. 2026). Additionally, improvements in the description of the spin dynamics during the inspiral were incorporated in IMRPhenomXPHM_SpinTaylor (Colleoni et al. 2025a), performing numerically the integration of the PN spin dynamic equations and enhancing the accuracy for the inspiral stage of the signal. Similar methods to improve the numerical integration were developed, implemented and published independently in Yu et al. (2023) and Gamba et al. (2022). The improvements introduced in IMRPhenomXO4a (including the dominant multipole equatorial asymmetry (Ghosh et al. 2024)) and in IMRPhenomXPHM_SpinTaylor, relative to IMRPhenomXPHM, have now been combined into a single model, IMRPhenomXPNR (Hamilton et al. 2026).

The SEOBNR approach has recently been developed with a new and more accurate generation of BBH waveform models, SEOBNRv5HM (Pompili et al. 2023) for aligned-spin systems and SEOBNRv5PHM (Ramos-Buades et al. 2023) for spin-precessing systems. Improvements include more analytical PN information (Henry 2023; Khalil et al. 2023) and

recent developments from second-order gravitational self-force for the flux and gravitational modes (Warburton et al. 2021; van de Meent et al. 2023), increasing the coverage of parameter space employed in the calibration with NR and test-mass limit simulations (Barausse et al. 2012; Taracchini et al. 2014a). Equatorial asymmetry effects for spin-precessing binaries have been incorporated as described in (Estellés et al. 2026). There are also substantial improvements to the computational efficiency of the models via the implementation of the post-adiabatic technique (Nagar & Rettegno 2019; Mihaylov et al. 2021) and the release of the modular and highly optimized Python package pySEOBNR (Mihaylov et al. 2023). The dominant and multipolar spin-aligned models SEOBNRv5 and SEOBNRv5HM also provide efficient ROM versions in the Fourier domain, SEOBNRv5_ROM and SEOBNRv5HM_ROM (Pompili et al. 2023).

2.2. *BNS Models*

When modeling BNS mergers, the inclusion of tidal interactions is critical for accurately describing the two-body dynamics and the corresponding GW emission. During the inspiral, these interactions are primarily characterized by the dimensionless tidal deformability parameter Λ, which encodes the response of a NS to the gravitational field of its companion and is directly related to the equation of state (EoS) of dense nuclear matter (Flanagan & Hinderer 2008; Hinderer 2008). Measurements of the tidal deformability through GW observations can constrain the EoS of dense nuclear matter (Chatziioannou 2020). Tidal interactions in BNS systems are first captured at 5 PN order (Flanagan & Hinderer 2008; Vines et al. 2011) in the velocity expansion and have been extended up to 7.5 PN order (Damour et al. 2012; Henry et al. 2020). Tidal corrections have been added to the GW phase of the PN inspiral Taylor Fourier-domain model TaylorF2 (Vines et al. 2011).

Besides being tidally deformed by their companions, rapidly-spinning NSs acquire an intrinsic oblateness where the NS mass-quadrupole moment depends on the dimensionless spin and encodes information about the EoS. The interaction between this quadrupole and the monopole of the companion generates an additional conservative potential, usually termed quadrupole–monopole coupling, entering the two-body dynamics at the same 2 PN order as the spin–spin term and contributing corrections to both the binding energy and the GW flux (Poisson 1998). Neglecting the quadrupole-monopole term can bias tidal-deformability and spin measurements once spins approach millisecond-pulsar values (Agathos et al. 2015; Samajdar & Dietrich 2018).

In addition to PN-based tidal corrections and quadrupole-monopole interaction, NR simulations play an essential role in accurately modeling the merger and post-merger phases of BNS coalescence, where nonlinear effects become significant, and several catalogs of BNS simulations have been produced (Dietrich et al. 2018; Gonzalez et al. 2023b; Kiuchi et al. 2017). The NRTidal approach (Dietrich et al. 2017) has been developed to incorporate tidal effects into the frequency-domain BBH waveform models, maintaining the computational efficiency of the BBH model baseline. The tidal phase in the frequency domain is modeled by employing the PN-expanded phase and calibrating it to BNS NR simulations. The original NRTidal model included PN-phase corrections augmented with calibration in the time domain using a non-spinning set of equal-mass BNS NR simulations, and then transformed to frequency-domain using the SPA. It was incorporated (Dietrich et al. 2019b) in the construction of the BNS models IMRPhenomD_NRTidal and SEOBNRv4_ROM_NRTidal, for aligned-spin systems, IMRPhenomPv2_NRTidal, for precessing systems, which was employed in the analysis of GW170817 (Abbott et al. 2017a) together with SEOBNRv4_ROM_NRTidal.

NRTidalv2 (Dietrich et al. 2019a) improved upon the previous model by the employment of improved NR data, the incorporation of spin effects, and the addition of amplitude corrections to the GW signal. It was incorporated to several BBH baselines models, including IMRPhenomPv2_NRTidalv2 (Dietrich et al. 2019a) and more recently IMRPhenomXP_NRTidalv2 (Colleoni et al. 2025b), for precessing systems, and SEOBNRv4_ROM_NRTidalv2 for aligned-spin systems.

Within the EOB framework, the SEOBNRv4T (Hinderer et al. 2016; Steinhoff et al. 2016) and TEOBResumS (Bernuzzi et al. 2015; Nagar et al. 2018; Akcay et al. 2019) models aim to provide a more accurate description of the inspiral phase, particularly in regimes where tidal interactions become significant. SEOBNRv4T includes tidally-induced multipole moments (up to $\ell = 3$), the effect of dynamical tides from f-mode resonances (Hinderer et al. 2016; Steinhoff et al. 2016), and the spin-induced quadrupole moment (Poisson 1998; Harry & Hinderer 2018). An efficient surrogate version of the model in the frequency domain, based on Gaussian process regression, was developed to reduce the cost of its employment in data-analysis applications (Lackey et al. 2019).

Similarly, TEOBResumS incorporates tidally-induced quadrupole moments using a different resummation and new information from gravitational self-force, with the lack of f-mode resonances but the addition of calibration to NR BNS simulations in a more recent version of the model (Gamba et al. 2023b). In order to improve its efficiency in data-analysis applications, a combination of the post-adiabatic technique and the SPA is employed to produce fast frequency-domain inspiral templates (Nagar & Rettegno 2019; Gamba et al. 2021).

2.3. *NSBH Models*

Modeling NSBH mergers presents distinct challenges compared to BNS systems, primarily because the NS may be fully disrupted and accreted by the BH without producing a post-merger remnant. As in BNS systems, the tidal response of the NS affects the phase evolution of the inspiral, requiring the inclusion of tidal corrections in the waveform phase (Pannarale et al. 2011). Additionally, in an EoS-dependent region of parameter space, typically involving unequal masses or

highly-spinning BHs, the NS can be tidally disrupted before reaching the innermost stable circular orbit. In such cases, the signal amplitude is strongly suppressed at high frequencies (Kyutoku et al. 2011; Foucart et al. 2014; Kawaguchi et al. 2015). While tidal disruption and post-merger remnants can also occur in BNS mergers, disrupted NSBH systems may lack a distinct merger signature altogether if the NS is fully disrupted before crossing the BH's horizon.

Current NSBH models employed to analyse the NSBH candidate signals from GWTC-2.0 onwards are the SEOBNRv4_ROM_NRTidalv2_NSBH (Matas et al. 2020) model and IMRPhenomNSBH (Thompson et al. 2020) model. Both models incorporate tidal information into the phase of the respective Fourier-domain BBH baseline model using the NRTidalv2 phase model described in Section 2.2. Disruptive and non-disruptive mergers are handled by adding tidal correction to the GW amplitude, as well as establishing a parameter-space dependent cut-off frequency for suppressing the GW amplitude in the case of disruptive events. In particular, IMRPhenomNSBH employs an NR-calibrated amplitude model (Pannarale et al. 2013, 2015), while SEOBNRv4_ROM_NRTidalv2_NSBH implements NR-calibrated correction factors to the underlying BNS NRTidal model. Both the amplitude model of SEOBNRv4_ROM_NRTidalv2_NSBH and the correction factors in SEOBNRv4_ROM_NRTidalv2_NSBH incorporate disruptive events by establishing a cut-off frequency and tapering the waveform at frequencies above this cut-off.

One of the main limitations of current NSBH models is that they are restricted to the dominant multipole and assume spin-aligned configurations. Including only the dominant mode can lead to degeneracies in distance and inclination, as well as to a non-negligible loss in signal-to-noise ratio (SNR) for unequal-mass NSBH systems, since higher-order multipoles can contain a significant fraction of the signal power for asymmetric binaries. Therefore, future models will need to address this limitation, and there is active development towards this (Gonzalez et al. 2023a).

2.4. *Luminosity and Remnant Properties*

Besides modeling the GW waveform, accurate predictions of the mass and spin of the final remnant BH, as well as estimations of the peak luminosity, are also crucial for population studies and tests of general relativity (GR). Several analytical fits have been developed using the remnant properties from NR BBH simulations (Hofmann et al. 2016; Keitel et al. 2017; Healy & Lousto 2017; Jiménez-Forteza et al. 2017), and recently accurate Gaussian process regression surrogate models have been developed (Varma et al. 2019; Islam et al. 2023). For NSBH systems, dependency on the tidal deformability of the NS have been included (Zappa et al. 2019). These fitting formulae generally depend on the component masses and spins, typically specified at some reference frequency in the inspiral stage of the signal. For estimating these quantities from the inferred source properties (see Section 5.9), the input spin values are evolved forward until a fiducial orbital frequency near merger is achieved, and then results from different fitting formulae are averaged to provide the final estimates. For precessing systems, these formulae are often augmented with in-plane spin contribution before the average procedure is applied (Johnson-McDaniel et al. 2016).

3. SIGNAL IDENTIFICATION

The GW strain time series produced by advanced-era interferometric detectors can be treated as a linear superposition of continuous non-astrophysical *noise* and occasional transient astrophysical *signals*. On timescales of tens of seconds, the noise is generally well approximated by colored stationary Gaussian processes, allowing for relatively stable modeling and analysis (Abbott et al. 2020a). However, on longer timescales, the noise becomes non-stationary, exhibiting time-dependent statistical properties (Zackay et al. 2021; Mozzon et al. 2020; Mozzon et al. 2022; Kumar et al. 2025). It is frequently contaminated by transient non-Gaussian artifacts, known as *glitches* (Nuttall 2018; Glanzer et al. 2023; Soni et al. 2025), and is also affected by slowly time-varying broadband disturbances (Abbott et al. 2020a). On the other hand, the signals in the data are presently relatively rare, occurring at a rate of just a few per week. GW *searches* are thus required to perform statistical data reduction, taking in the kilohertz-sampled strain data and producing a list of astrophysical candidates. The candidate lists are released as part of the GWTC-5.0 data release (LIGO–Virgo–KAGRA Collaboration 2025a) and also used as input to the downstream data-analysis workflow illustrated in Figure 1.

The search for GW transient candidates is carried out through two distinct phases. Initially, *online* analyses enable prompt follow-up observations by the global astronomical community (Abbott et al. 2019b, 2023a; LIGO–Virgo–KAGRA Collaboration 2025b). Later, *offline* analyses are conducted to produce a more accurate list of candidates by re-evaluating the significance of the initial online candidates and identifying new ones. Some of the factors responsible for the differences in the online and offline analyses are the use of the final strain data, improved data quality, enhanced algorithms, and a better understanding of the noise background in the data.

Two broad families of detection algorithms are used: search pipelines for minimally-modeled transient sources that do not rely on specific waveform predictions and search pipelines that, conversely, aim to maximize sensitivity to CBCs by using sets of template waveforms (a subset of the models described in Section 2). For the GWTC, we use multiple pipelines concurrently and combine their results. This approach minimizes the risk of missing astrophysical signals and allows us to cover a wider range of the CBC parameter space.

Template-based pipelines target specific regions of the CBCs parameter space, usually defined by the redshifted (detector-frame) masses (Krolak & Schutz 1987) and spins of the two compact objects, which are arranged into *template banks*. Template waveforms are currently calculated neglect-

Table 1. A summary of waveform models used in each release of the GWTC. Since the catalog is cumulative, later releases (e.g., GWTC-5.0) include all previous candidates. This includes the re-analysis of O4a which we label GWTC-4.1. In most cases, PE results from earlier releases remain unchanged. An exception is GWTC-2.1, which reanalyzed O1 and O2 data using updated waveform models, replacing the original PE results from GWTC-1.0 and GWTC-2.0. As a result, the PE analyses with the waveform approximants listed for GWTC-1.0 and GWTC-2.0 are not part of GWTC-5.0. The models employed in GWTC-5.0 are discussed in Section 2, and their use within specific pipelines is described in the following sections. Special candidate analyses might employ additional waveform models not listed here, but discussed for those particular signals.

Catalog release	Data analysed	Search templates	Sensitivity estimates	Parameter estimation
GWTC-1.0 (Abbott et al. 2019a)	O1, O2	SEOBNRv4_ROM, TaylorF2	-	IMRPhenomPv2, IMRPhenomPv2_NRTidal, SEOBNRv3, SEOBNRv4_ROM_NRTidal, SEOBNRv4T, TaylorF2, TEOBResumS
GWTC-2.0 (Abbott et al. 2021a)	O3a	SEOBNRv4_ROM, TaylorF2	SEOBNRv4_opt	IMRPhenomD, IMRPhenomD_NRTidal, IMRPhenomHM, IMRPhenomPv2, IMRPhenomPv2_NRTidal, IMRPhenomPv3HM, IMRPhenomNSBH, NRSur7dq4, SEOBNRv4_ROM_NRTidalv2_NSBH, SEOBNRv4_ROM, SEOBNRv4HM_ROM, SEOBNRv4P, SEOBNRv4PHM, SEOBNRv4T_surrogate, TaylorF2, TEOBResumS
GWTC-2.1 (Abbott et al. 2024)	O1–O3a	SEOBNRv4_opt, SEOBNRv4_ROM, SpinTaylorT4, TaylorF2	SEOBNRv4P, SEOBNRv4PHM, SpinTaylorT4	IMRPhenomPv2_NRTidal, IMRPhenomXPHM, SEOBNRv4PHM,
GWTC-3.0 (Abbott et al. 2023a)	O3b	SEOBNRv4_opt, SEOBNRv4_ROM, SpinTaylorT4, TaylorF2	SEOBNRv4P, SEOBNRv4PHM, SpinTaylorT4	IMRPhenomNSBH, IMRPhenomXPHM, SEOBNRv4_ROM_NRTidalv2_NSBH, SEOBNRv4PHM
GWTC-4.0 (Abac et al. 2025a)	O4a	IMRPhenomD, SEOBNRv4_opt, SEOBNRv4_ROM, SEOBNRv5_ROM, SpinTaylorT4, TaylorF2	IMRPhenomXPHM	IMRPhenomNSBH, IMRPhenomPv2_NRTidalv2, IMRPhenomXO4a IMRPhenomXPHM_SpinTaylor, NRSur7dq4, SEOBNRv4_ROM_NRTidalv2_NSBH, SEOBNRv5PHM
GWTC-5.0 (Abac et al. 2026b)	O4b (and O4a)	IMRPhenomD, SEOBNRv4_opt, SEOBNRv4_ROM, SEOBNRv5_ROM, SpinTaylorT4, TaylorF2	IMRPhenomXPHM	IMRPhenomXPHM_SpinTaylor, IMRPhenomXPNR, NRSur7dq4, SEOBNRv5PHM

ing several physical effects, specifically non-dominant multipole emission, orbital precession, eccentricity and tidal deformability. The templates thus only model the dominant multipole emission from quasi-circular BBH coalescences with spins parallel to the orbital angular momentum.

This simplification, which is appropriate for the majority of the CBC signals detected so far, reduces the dimensionality of the parameter space to be covered, leading to a beneficial reduction of the computational and operational costs of the analysis. The price to pay for this simplification is a reduced search sensitivity in regions of the CBC parameter space where the neglected physical effects give a large contribution to the signal (Dal Canton et al. 2015; Calderón Bustillo et al. 2016; Harry et al. 2016; Calderón Bustillo et al. 2017; Dietrich et al. 2019b; Ramos-Buades et al. 2020b; Phukon et al. 2025; Chia et al. 2024; Mehta et al. 2025). Combining template-based with minimally-modeled search methods partially covers such cases, since the latter impose only weak assumptions on the signal. Moreover, signals that appear to deviate from the simplified templates, such as GW190412 (Abbott et al. 2020c) and GW190521 (Abbott et al. 2020d), can still be detected with high significance because the deviation is relatively small. Still, including the correct physical effects in a template-based search is feasible and has yielded additional marginally-significant candidate events with interesting properties (Dhurkunde & Nitz 2025; Wadekar et al. 2024).

In modeled searches, we search the data from each detector by matched filtering each normalized template $\tilde{u}(f|\boldsymbol{\theta}_{\rm int})$ (where $\boldsymbol{\theta}_{\rm int}$ is a vector over a subset of intrinsic CBC parameters), to produce a SNR time series defined as (Maggiore 2007; Allen et al. 2012)

$$\rho(t) = 4\left|\int_{f_{\rm low}}^{f_{\rm high}} \frac{\tilde{d}(f)\tilde{u}^*(f|\boldsymbol{\theta}_{\rm int})}{S_n(f)} e^{2i\pi ft} df\right|, \qquad (1)$$

Table 2. Waveform models used in GWTC analyses. We indicate if the model includes spin-precession, matter effects and which multipoles are included for each model. For spin-precessing models, the multipoles correspond to those available in the coprecessing frame. We indicate the most relevant reference for each model, additional references are given in the main text of Section 2. Special candidate analyses might employ additional waveform models not listed here, but discussed for those particular signals.

Model	Precession	Multipoles $(\ell, \lvert m\rvert)$	Matter	Reference
SpinTaylorT4	✓	$\ell \leq 4$	✓	Isoyama et al. (2020)
TaylorF2	×	$(2,2)$	✓	Damour et al. (2001)
IMRPhenomD	×	$(2,2)$	×	Khan et al. (2016)
IMRPhenomD_NRTidal	×	$(2,2)$	✓	Dietrich et al. (2019b)
IMRPhenomHM	×	$(2,2),(2,1),(3,3),(3,2),(4,4),(4,3)$	×	London et al. (2018)
IMRPhenomNSBH	×	$(2,2)$	✓	Thompson et al. (2020)
IMRPhenomPv2	✓	$(2,2)$	×	Bohé et al. (2016)
IMRPhenomPv2_NRTidal	✓	$(2,2)$	✓	Dietrich et al. (2019b)
IMRPhenomPv2_NRTidalv2	✓	$(2,2)$	✓	Dietrich et al. (2019a)
IMRPhenomPv3HM	✓	$(2,2),(2,1),(3,3),(3,2),(4,4),(4,3)$	×	Khan et al. (2020)
IMRPhenomXPHM	✓	$(2,2),(2,1),(3,3),(3,2),(4,4)$	×	Pratten et al. (2021)
IMRPhenomXPHM_SpinTaylor	✓	$(2,2),(2,1),(3,3),(3,2),(4,4)$	×	Colleoni et al. (2025a)
IMRPhenomXO4a	✓	$(2,2),(2,1),(3,3),(3,2),(4,4)$	×	Thompson et al. (2024)
IMRPhenomXPNR	✓	$(2,2),(2,1),(3,3),(3,2),(4,4)$	×	Hamilton et al. (2026)
NRSur7dq4	✓	$\ell \leq 4$	×	Varma et al. (2019)
SEOBNRv3	✓	$(2,2),(2,1)$	×	Pan et al. (2014)
SEOBNRv4_opt	×	$(2,2)$	×	Devine et al. (2016)
SEOBNRv4_ROM	×	$(2,2)$	×	Bohé et al. (2017)
SEOBNRv4HM_ROM	×	$(2,2),(2,1),(3,3),(4,4),(5,5)$	×	Cotesta et al. (2020)
SEOBNRv4_ROM_NRTidal	×	$(2,2)$	✓	Dietrich et al. (2019b)
SEOBNRv4_ROM_NRTidalv2_NSBH	×	$(2,2)$	✓	Matas et al. (2020)
SEOBNRv4P	✓	$(2,2),(2,1)$	×	Ossokine et al. (2020)
SEOBNRv4PHM	✓	$(2,2),(2,1),(3,3),(4,4),(5,5)$	×	Ossokine et al. (2020)
SEOBNRv4T	×	$(2,2)$	✓	Steinhoff et al. (2016)
SEOBNRv4T_surrogate	×	$(2,2)$	✓	Lackey et al. (2019)
SEOBNRv5_ROM	×	$(2,2)$	×	Pompili et al. (2023)
SEOBNRv5PHM	✓	$(2,2),(2,1),(3,3),(3,2),(4,4),(4,3),(5,5)$	×	Ramos-Buades et al. (2023) Estellés et al. (2026)
TEOBResumS	×	$(2,2),(2,1),(3,3),(3,2),(3,1)$	✓	Akcay et al. (2019)

where $\tilde{d}$ represents the strain data, S_n is the power spectral density (PSD) modeling the detector noise (see Appendix B of Abac et al. 2026a), and the tilde over a symbol indicates the Fourier transform of a time series.

The lower integration limit frequency $f_{\rm low}$ is determined by considerations of computational cost, data quality, and calibration reliability. The upper limit $f_{\rm high}$ is predominantly based on considerations of computational cost. Both limits may vary with search implementations and templates. In practice, S_n is a point estimate of the noise variance at a given time and frequency, and is measured empirically from the data themselves. This is possible in a straightforward way with current detectors thanks to the sparsity of the transient astrophysical signals and the absence of detectable persistent signals. This measurement is repeated at regular intervals in order to track the time-dependence of noise properties.

Modeled search pipelines produce triggers by looking for peaks exceeding a given threshold in the SNR time series. This same procedure is applied to the calibrated data provided by all available detectors.

All search pipelines, minimally-modeled or templated-based, attempt to pair candidates from one detector with coincident candidates from others, though some pipelines also consider triggers observed in a single interferometer (Sachdev et al. 2019; Alléné et al. 2025; Cabourn Davies & Harry 2022). Some pipelines employ an additional *coherent* step, where the signal likelihood is maximized over some of the extrinsic parameters using strain or SNR time series information from all detectors simultaneously, rather than only from SNR maxima in each separate detector. The root-sum-square of the single-detector SNRs, called the *network SNR*, is often used to quantify the overall loudness of a candidate over all available detectors (Cutler & Flanagan 1994; Maggiore 2007). Each pipeline combines the SNR information with additional signal-consistency tests to construct its own ranking statistic.

Based on their ranking, search pipelines assign each candidate a false-alarm rate (FAR) value, which quantifies the expected rate of non-astrophysical candidates produced with a rank at least as high as the candidate under consideration.

In addition to the measure of detection significance provided by the FAR of a candidate, we have also provided the probability $p_{\rm astro}$ that a candidate is of astrophysical rather than terrestrial (noise) origin. The calculation of $p_{\rm astro}$ accounts for both the rate of noise events corresponding to the FAR and an estimated rate of astrophysical signals under some prior model of the CBC population (Farr et al. 2015; Kapadia et al. 2020). It is formally defined as (Abbott et al. 2016a)

$$p_{\rm astro}(x) = \frac{\mathcal{R}_{\rm astro} f(x)}{\mathcal{R}_{\rm astro} f(x) + \mathcal{R}_{\rm noise} b(x)}, \tag{2}$$

where $\mathcal{R}_{\rm astro}$ is the total rate of astrophysical candidates in a given pipeline, $\mathcal{R}_{\rm noise}$ is the total rate of candidates due to noise; $f(x)$ is the probability density function (pdf) of signal events at the candidate ranking value x, and $b(x)$ is the equivalent pdf for noise events. Each pipeline estimates these rates and calculates a $p_{\rm astro}$ value for the candidates it identifies using its own method.

Since O4a, most pipelines use a common signal model for $f(x)$ based on a distribution over source masses, spins and distance resembling the astrophysical population inferred from GWTC-3.0 (Abbott et al. 2023b) while supporting possible outliers (Essick et al. 2025); a slightly different model used by one analysis is described in Section 3.2.

The total astrophysical probability $p_{\rm astro}$ is then distributed between three mutually-exclusive source categories: a BNS class corresponding to both component (source-frame) masses ranging between $1\,M_\odot$ and $3\,M_\odot$; a NSBH class for one mass between $1\,M_\odot$ and $3\,M_\odot$ and the other above; and a BBH class corresponding to both masses above $3\,M_\odot$. The associated probabilities for each category thus sum to $p_{\rm astro}$:

$$p_{\rm BNS} + p_{\rm NSBH} + p_{\rm BBH} = p_{\rm astro}. \tag{3}$$

The complementary probability that the candidate is of terrestrial origin, i.e., caused by noise rather than a CBC, is notated as $p_{\rm terr} = 1 - p_{\rm astro}$.

Transient search methods used to compile our catalog for data from the first observing run (O1) and second observing run (O2) are described in Section 3 of Abbott et al. (2019a), and for data from third observing run (O3), in Section 4 of Abbott et al. (2021a), Section 3 of Abbott et al. (2024) and Section 4 of Abbott et al. (2023a). We describe here the methods used for data from O4a and O4b, in the context of this cumulative catalog. In the coming subsections, we present details of the different pipelines used to identify transient GW candidates, including coherent WaveBurst (cWB) for minimally-modeled sources (Klimenko et al. 2016; Mishra et al. 2022, 2025), and GStreamer LIGO Algorithm Library (GstLAL; Cannon et al. 2012a; Messick et al. 2017; Sachdev et al. 2019; Hanna et al. 2020; Cannon et al. 2021; Sakon et al. 2024; Ray et al. 2023; Tsukada et al. 2023; Ewing et al. 2024; Joshi et al. 2025a,b), Multi-Band Template Analysis (MBTA; Alléné et al. 2025), PyCBC (Dal Canton et al. 2021), and Summed Parallel Infinite Impulse Response (SPIIR; Chu et al. 2022) for model-based search analyses. The names we use for the different pipelines typically indicate publicly available and fairly general software packages. In principle, one may construct many different search algorithms with those packages, and may configure them in arbitrary ways to optimize different figures of merit (e.g., sensitivity to specific populations of signals, or computing cost). For the purposes of GWTC, each pipeline name identifies the combination of a specific software package with the specific set of configuration choices described in this paper. In Table 3, we summarize the key features of each search pipeline and the strain data they analyzed. The table focuses on the configurations used for offline analyses, which are the main scope of this paper, though the online configurations are similar.

Most of these search pipelines also participated in pre-GWTC-4.0 releases, with some configuration differences. The main change concerning modeled searches is that some

pipelines have refined their search space to reflect past detections (Abbott et al. 2023a) and to cover the higher redshift range enabled by improved detector sensitivity. In addition, all pipelines have implemented new features aimed at better distinguishing non-astrophysical transients and improving sensitivity to CBC signals. In parallel, they continue to improve their computational and operational efficiency. There is also a continuous effort to improve the population prior used in searches, informed by results on past observing runs (Abbott et al. 2023b). The main differences for each pipeline are described in the corresponding subsections below.

Each search pipeline provides its own list of candidates, along with corresponding values of the FAR and $p_{\rm astro}$, which are neither corrected for multiple trials nor combined in any way (see Section 3.6).

For offline analyses, a set of *Category 1* (CAT1) flags are determined, which flag periods of severe data-quality or technical issues affecting the interferometers during the observation time. During CAT1 periods (detailed in Section 4.1 of Soni et al. 2025), the data are considered too contaminated by noise to be analyzable within realistic computational or human resources.

During the O1 to O3 observing runs, short-duration *Category 2* (CAT2) flags were also provided for transient searches (Abbott et al. 2020a). CAT2 flags are based on statistical correlations found between excess noise transients in auxiliary channels and in the strain data. As of O4, these CAT2 flags are no longer used in CBC searches, though they continue to be used for minimally-modeled transient searches (Soni et al. 2025).

As an extra input for CBC searches, a time series is generated by the IDQ pipeline (Essick et al. 2020). This time series contains statistical data-quality information based on the activity of auxiliary channels deemed safe (i.e., channels that should not be influenced by GWs). Observation times marked with IDQ flags are likely to be contaminated by glitches. Unlike CAT1 flags, which are mandatory vetoes prior to any analysis, IDQ flags are optional and may be used either as vetoes, or as input in the ranking statistic of the candidates. CAT1 vetoes and the IDQ time series are created from the strain data and auxiliary channels as a preprocessing step and ingested by the offline searches (see Figure 1). Details of the data release of the IDQ time series are provided in LIGO–Virgo–KAGRA Collaboration (2025c).

3.1. CWB-BBH

The CWB algorithm is designed to detect transient GW signals using networks of GW detectors without requiring a specific waveform model.

It identifies coincident excess power events by analyzing multi-resolution time–frequency (TF) representations of detector data (Klimenko et al. 2008, 2016) whitened in the Wilson–Daubechies–Meyer (WDM) wavelet domain (Necula et al. 2012).

Once potential GW signals or noise events are identified, CWB reconstructs the source sky location and the signal waveforms recorded by the detectors using the constrained maximum-likelihood method (Klimenko et al. 2005, 2016).

This analysis is performed on detector strain data, and can in principle accommodate signal frequencies up to a few kilohertz and durations up to a few seconds, though one might choose a narrower search space as part of the configuration choices (see Table 3).

The CWB detection statistic relies on coherent energy $E_{\rm c}$, which is derived from cross-correlating the reconstructed signal waveforms in the detectors and normalizing them by the spectral amplitude of the detector noise. The square root of $E_{\rm c}$ provides a lower bound on the signal network SNR and is used for the initial selection of candidates identified by CWB. As the production rate of CWB candidates is primarily driven by glitches, additional post-production vetoes are applied to further reduce the background rates.

One primary signal-independent veto statistics is the network correlation coefficient, defined as $c_{\rm c} = E_{\rm c}/(E_{\rm c} + E_{\rm n})$, where $E_{\rm n}$ represents the normalized residual noise energy after the reconstructed signal has been subtracted from the data. Typically, for glitches $c_{\rm c} \ll 1$. Consequently, candidates with $c_{\rm c} < 0.6$ are rejected as potential glitches. For authentic GW signals, the residual energy $E_{\rm n}$ is expected to follow the χ^2 distribution with the number of degrees of freedom $N_{\rm DoF}$ proportional to the number of TF data samples comprising the signal. The reduced $\tilde{\chi}^2 = \chi^2/N_{\rm DoF}$ statistic serves as a powerful signal-independent veto, effectively identifying and removing glitches where $\tilde{\chi}^2$ is significantly greater than 1. To enhance the search for CBC signals, CWB performs analysis in the frequency band below 512 Hz and also employs weak signal-dependent vetoes. These vetoes are based on the central frequency $f_{\rm c}$ and the detector-frame chirp mass $(1+z)\mathcal{M}$, both of which are estimated from the TF evolution of the signal power (Tiwari et al. 2016).

All candidates identified by CWB are ranked by the reduced coherent network SNR

$$\rho_{\rm cWB} = \sqrt{\frac{E_{\rm c}}{1 + \tilde{\chi}^2[\max(1, \tilde{\chi}^2) - 1]}}\,. \tag{4}$$

To estimate the statistical significance of the GW candidates, they are ranked against background events generated by the detector noise. A comprehensive background data sample, equivalent to approximately 1000 years of observation time, is obtained by repeating the CWB analysis on time-shifted data. To ensure that astrophysical signals are excluded from the background data, the time shifts are selected to be much larger than the expected signal time delay between the detectors. The statistical significance of each CWB candidate is quantified by its FAR. The FAR is defined as the rate of background events that exhibit a larger $\rho_{\rm cWB}$ value than the GW candidate in question. To account for potential long-term variations in detector noise, the candidate significance is estimated using nearby data intervals, typically one to two weeks in length. For the calculation of $p_{\rm astro}$, CWB combines the

Table 3. A high-level overview of the similarities and differences between GW search pipelines used in the analysis of O4b data for GWTC-5.0. The `GDS-CALIB_STRAIN_CLEAN` and `GDS-CALIB_STRAIN_CLEAN_AR` channels are identical other than that the latter contains only data for which the detectors are nominally in a good working state (Abac et al. 2026c).

	cWB-BBH	GstLAL	MBTA	PyCBC	SPIIR (online only)
Data source (online) channel name	`GDS-CALIB_STRAIN_CLEAN` (LIGO)	`GDS-CALIB_STRAIN_CLEAN` (LIGO) `Hrec_hoft_16384Hz` (Virgo)			
Data source (offline) channel name	`GDS-CALIB_STRAIN_CLEAN` (LIGO)	`GDS-CALIB_STRAIN_CLEAN` (LIGO) `Hrec_hoft_16384Hz` (Virgo)	`GDS-CALIB_STRAIN_CLEAN_AR` (LIGO) `Hrec_hoftReproc1AR_16384Hz` (Virgo)		N/A N/A
Target waveform morphology	Linear superposition of wavelets	Quasicircular spin-aligned $(2,2)$-mode CBCs			
PSD estimation	WDM wavelet	Welch with geometric-mean and median averaging	Welch with median averaging		
Filtering approach	Time-frequency decomposition	Matched filtering via time-domain SVD filters	Matched filtering via multi-band FFT	Matched filtering via single-band FFT	Matched filtering via time-domain summed parallel IIR filters
Working frequency band (Hz)	16–320	10–1024 (template dependent)	20–2048 (template dependent)	15–1024 (template dependent)	15–1024
Glitch suppression	Gating, coherence test	Mass-based gating, autocorrelation ξ^2 test	Gating, autocorrelation χ^2 test, trigger-rate penalty	Gating, time-frequency χ^2 test, sine–Gaussian χ^2 test	Gating, autocorrelation χ^2 test
Multi-detector trigger formation	Coherent	Coincident			Coherent
Single-detector candidates	N/A	IMBH excluded; Other singles penalized	$\mathcal{M} < 7M_\odot$	Template duration > 0.3 s, background-based cuts	N/A
Ranking statistic	Coherent network SNR weighted by ML classifier	Likelihood ratio	Reweighted SNR with astrophysical prior	Likelihood ratio with KDE-based template weighting	Likelihood ratio, extrapolation using KNN KDE
Use of external data-quality info	CAT2 as veto for candidates	None		IDQ included in ranking	None
Multi-detector background estimation	Time shifts	Monte Carlo of single-detector trigger distributions, excluding manually-vetted online candidates	Time shifts		
Single-detector background estimation	N/A		Exponential fit and extrapolation of empirical distribution		N/A
p_{astro} population model	Mass and spin distributions based on GWTC-3.0 population (BBH only)	Salpeter primary mass function	Mass and spin distributions based on GWTC-3.0 population		

time-shifted background analysis and simulated BBH signals based on the GWTC-3.0 population (Abbott et al. 2023a). Only p_{BBH} is calculated.

Since the publication of GWTC-3.0, the cWB algorithm used in the previous observing runs (Klimenko et al. 2016; Drago et al. 2021), i.e., cWB-2G, has undergone significant changes, and a new version, cWB-BBH, has been adopted for CBC searches: it incorporates two major changes compared to cWB-2G.

The first change involves replacing the WDM wavelet transform (Necula et al. 2012) with the multi-resolution *WaveScan* transform, which is based on the Gabor wavelets (Klimenko 2022). This update effectively reduces temporal and spectral leakage in the TF data, potentially leading to more accurate signal representation.

Secondly, in addition to the traditional excess-power statistic, the algorithm now incorporates the cross-power statistic (Klimenko 2022) for identifying transient signals. The excess-power represents the total power of a TF data sample (or *WaveScan* pixel) integrated over the detector network. The cross-power represents the correlation of power in the detectors. Both statistics are maximized over possible time-of-flight delays of a GW signal across the detector network. The excess-power and the cross-power amplitudes, a_{e} and $a_{\times}$, respectively, follow a predictable half-normal distribution with unity variance assuming quasi-stationary detector noise (Klimenko 2022). For analysis, TF pixels with the excess-power amplitude exceeding 2.3 standard deviations are selected. Adjacent pixels in the TF plane are then clustered to form the initial cWB-BBH candidates (clusters).

The clustered excess-power $\sum_{i=1}^{I} a_{\mathrm{e}}^2[i]$ and the cross-power $\sum_{i=1}^{I} a_{\times}^2[i]$ are used to calculate the upper bounds on the signal network SNR ρ_{e} and the coherent network SNR $\rho_{\times}$ respectively (Mishra et al. 2025), where I is the number of the TF pixels comprising the cluster.

The ρ_{e} is optimized for identification of clusters due to fluctuations of the quasi-stationary detector noise dominating the initial cWB-BBH rate on the order of 100 Hz.

All clusters with $\rho_{\mathrm{e}} > 4$ are accepted for further analysis.

The accepted cWB-BBH clusters are mostly produced by glitches with a typical rate of 0.1 Hz. After selection, the initial TF clusters are aggregated if they fall within the time and frequency intervals of 0.23 s and 64 Hz respectively. This improves the energy collection for transient candidates that can be fragmented into clusters by the TF transform. The glitch rate is further reduced by requiring $\rho_{\times} > 7$, where the upper bound of the coherent network SNR is calculated for the defragmented candidates.

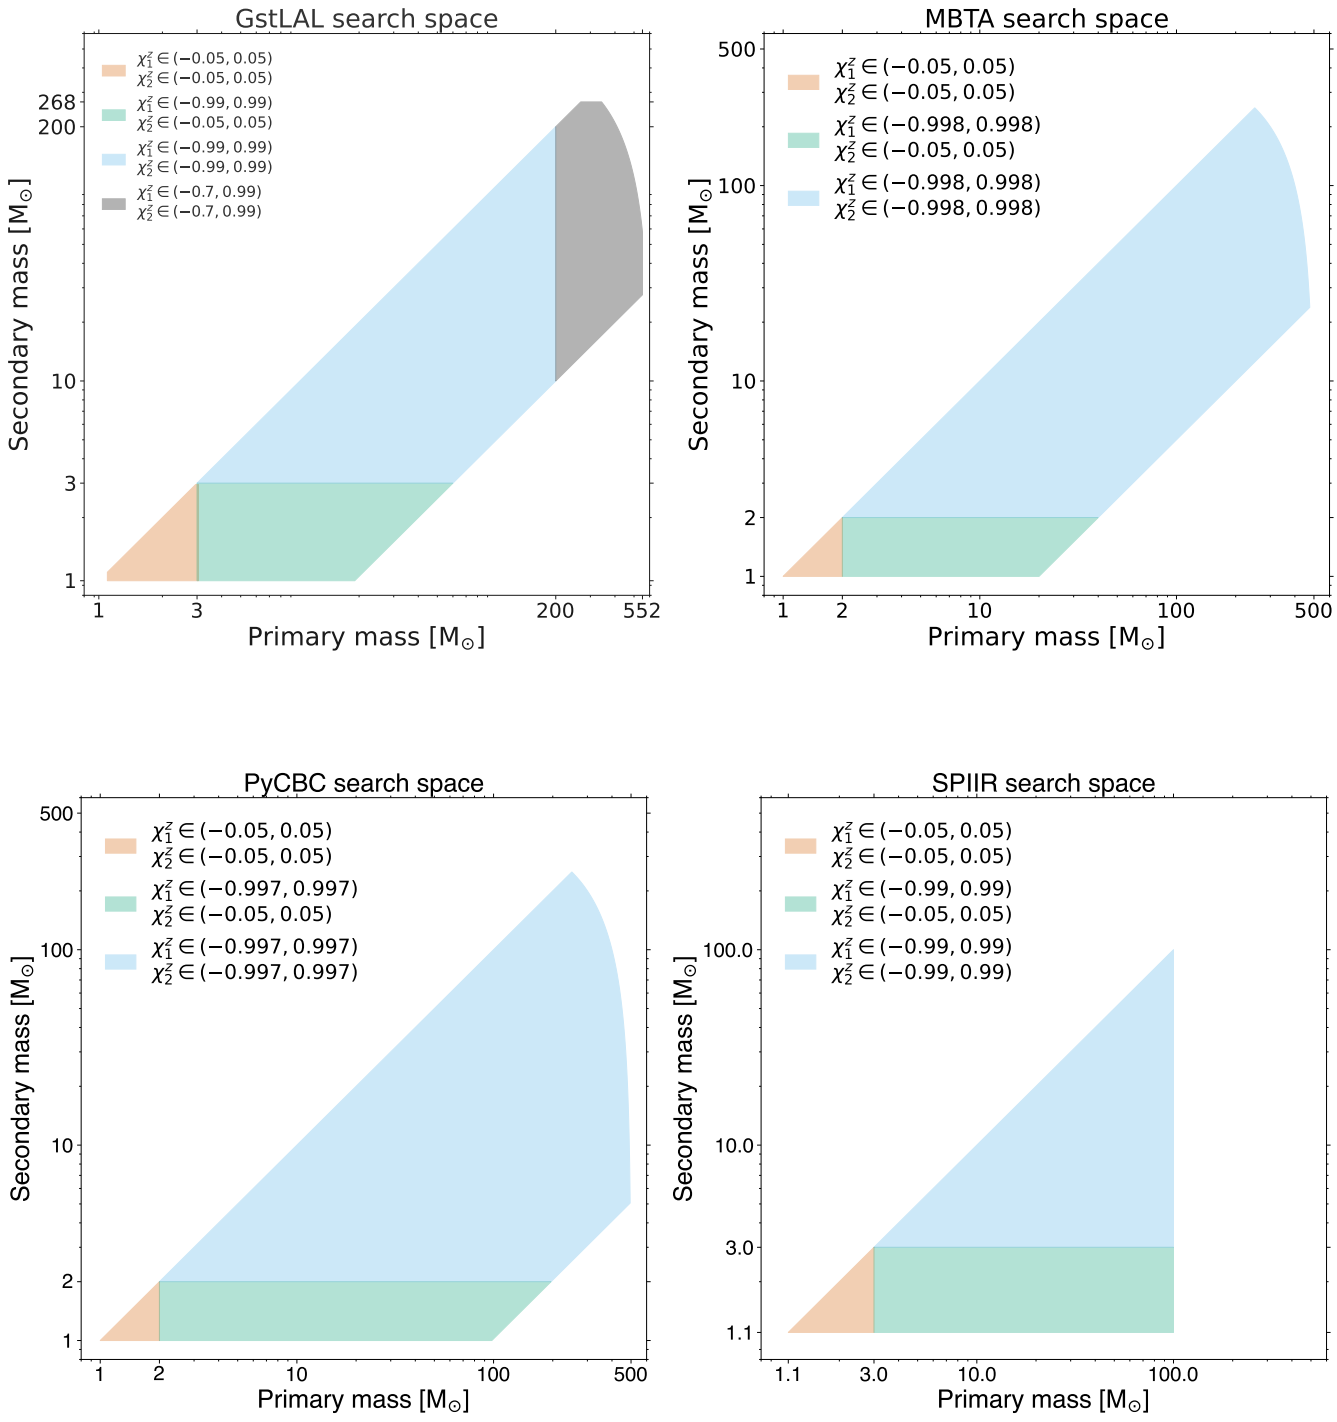

Figure 2. Regions of the CBC parameter space explored by our different template-based search pipelines during the O4b observing run. Individual masses are given in the detector frame, i.e. they are redshifted. The spin parameters $\chi^z_{1,2}$ are the spin projections along the orbital angular momentum.

The above conditions on the $\rho_{\rm e}$ and $\rho_{\times}$ require that the GW signal clusters with SNR > 4 and defragmented GW candidates with SNR > 7 are accepted for the analysis. The resulting reduced candidates rate, on the order of $1\,\mathrm{mHz}$, makes the subsequent likelihood analysis and reconstruction of the remaining candidates computationally feasible. At this stage, the sky location of each candidate is determined and the signal waveforms are reconstructed with the inverse *WaveScan* transform (Klimenko 2022). The cWB-BBH candidates with $\rho_{\rm cWB} > 7$ are stored for the post-production analysis.

Moreover, the post-production veto analysis has been replaced by the XGBoost machine-learning algorithm (Mishra et al. 2021; Mishra et al. 2022, 2025), based on an ensemble of decision trees. It performs a classification of cWB-BBH candidates by using a subset of 25 summary statistics including $\rho_{\rm cWB}, c_{\rm c}, \tilde{\chi}, \rho_{\rm e}, \rho_{\times}, f_{\rm c}$, and $(1+z)\mathcal{M}$. To improve the distinction between astrophysical GW signals and glitches, the cWB-BBH detection statistic is modified as

$$\rho_{\rm r} = \rho_{\rm cWB} W_{\rm XGB}, \tag{5}$$

where $W_{\rm XGB}$ is the XGBoost classification factor (Mishra et al. 2021) ranging from 0 (indicating a glitch) to 1 (indicating a signal). The XGBoost response is trained using a subset of background candidates and a representative set of simulated BBH signals, which are injected into the detector data and recovered with cWB-BBH. Due to a weak dependence of the cWB-BBH summary statistics on the CBC population and signal models, the algorithm is robust to a variety of CBC features, including higher multipoles, unequal component masses, misaligned spins, eccentric orbits, and possible deviations from GR (Mishra et al. 2022; Bhaumik et al. 2025).

While in principle the cWB algorithm can run with any number of interferometers, the version used in the O4b searches only runs with interferometer pairs. In particular, for this catalog cWB-BBH only searched data from the LIGO Hanford Observatory (LHO) and LIGO Livingston Observatory (LLO) interferometer pair.

3.2. GstLAL

GstLAL is a stream-based time-domain matched-filtering search pipeline (Cannon et al. 2012a; Messick et al. 2017; Sachdev et al. 2019; Hanna et al. 2020; Cannon et al. 2021; Sakon et al. 2024; Ray et al. 2023; Tsukada et al. 2023; Ewing et al. 2024; Joshi et al. 2025a,b). All previous GWTC versions included results produced by GstLAL (Abbott et al. 2019a, 2021a, 2024, 2023a; Abac et al. 2025a). For this analysis, multiple GstLAL configurations were used: an online configuration whose results were re-evaluated using the full background, an offline configuration that processed data dropped by the online configuration, and another offline configuration that targeted intermediate mass black hole (IMBH) mergers (Joshi et al. 2025a). The results published in GWTC-5.0, include analysis of data from the LIGO detectors during O4a, and from the LIGO and Virgo detectors during O4b. The primary configuration was an online search, which identified CBC signals in near real-time, enabling rapid alerts and follow-up. The online search dropped $\sim$2% of the observing run data due to, for example, computer downtime. The dropped data were processed in an offline configuration, methodologically identical to the online configuration, but executed post O4b. The online configuration did not cover the high-mass region of the parameter space, so an additional offline configuration was run to target the IMBH region, shown in grey in Figure 2. All configurations (online, dropped-data and IMBH) were used to evaluate search sensitivity by injecting simulated signals into the detector strain data and analyzing the effectiveness of signal recovery. Final candidate ranking and significance estimation incorporated the collective results from all configurations. An additional early-warning configuration, designed to identify BNS coalescences prior to merger, was also deployed during O4b. However, results from the early-warning configuration are not part of GWTC-5.0.

The analysis begins by constructing a template bank to cover the targeted astrophysical parameter space. Strain data are then pre-processed to prepare for matched-filtering, which correlates the pre-processed strain with templates from the bank. Peaks in the matched-filter output are identified as triggers and subsequently ranked using a likelihood ratio (LR) statistic that incorporates signal consistency and detector network properties (Tsukada et al. 2023). The background LR distribution from noise triggers is used to estimate the FAR and the significance of each GW candidate. In the following paragraphs, we describe each component of the analysis in detail, beginning with the template bank.

For O4b, the template bank targeted CBCs with detector-frame total masses between $1.95\,M_\odot$ and $610\,M_\odot$, and mass ratios from 0.05 to 1 (Sakon et al. 2024; Joshi et al. 2025a). It was divided into two disjoint regions: stellar-mass (component masses up to $200\,M_\odot$) and IMBH (primary component mass between $200\,M_\odot$ and $552\,M_\odot$ and secondary component mass below $268\,M_\odot$). While no BBHs with total mass above $400\,M_\odot$ have been observed so far, cosmological redshifting of source-frame masses can shift distant stellar-mass binaries into the IMBH region in the detector frame, motivating an extended search (Abbott et al. 2023a). Templates in both regions were spin aligned, i.e., spin components in the orbital plane are zero.

Aligned spin values were allowed in the range $[-0.99, +0.99]$ for stellar-mass templates, with additional restriction to $[-0.05, +0.05]$ for component masses below $3\,M_\odot$, consistent with observations for galactic BNSs (Stovall et al. 2018) and motivated by computational and background considerations.

In the IMBH region, spins were restricted to $[-0.70, +0.99]$ to mitigate triggers from noise transients with similar time–frequency structure as high-mass waveforms with large anti-aligned spins (Hanna et al. 2022). The template bank used is identical to the one used in O4a. The search space is illustrated in Figure 2.

During template bank creation, templates were placed using a metric based on the IMRPhenomD approximant

(Husa et al. 2016; Khan et al. 2016), with a lower frequency cutoff of 10 Hz and a maximum duration of 128 s, ensuring coverage of the detector's sensitive band while limiting computational cost. The duration constraint led to a template-dependent low-frequency cutoff for low-mass systems. For example, a $1.4\,M_\odot$–$1.4\,M_\odot$ binary would have a waveform duration of over 1000 s starting from 10 Hz. To stay within the 128 s constraint, its corresponding lower frequency cutoff must be raised to ∼22 Hz (Peters & Mathews 1963).

Template placement was performed using MANIFOLD (Hanna et al. 2023), an algorithm based on a geometric binary tree that tiles the parameter space dictated by the minimal match constraints. A minimal match of 97% was applied in the stellar-mass region, corresponding to a 10% expected loss in detection rate (Owen 1996), while the IMBH region used 99%, resulting in a total of $\sim 1.8 \times 10^6$ templates.

To improve fault tolerance of the online configuration, the stellar-mass bank was interleaved into two complementary halves, each independently covering the stellar-mass region and analyzed independently at separate data centers (Godwin 2020; Sakon et al. 2024). Each half was subdivided into template bins and each bin was decomposed into filters using singular value decomposition (SVD; Cannon et al. 2012b; Sakon et al. 2024).

After template bank construction, the detector strain data were preprocessed to ensure compatibility with the templates and to reduce noise contamination. Data were resampled to 2048 Hz and whitened. To suppress short-duration high-amplitude noise transients that can mimic astrophysical signals, especially at high masses where template waveforms are short, amplitude-based gating was applied to the whitened data. The gating threshold was a linear function of the chirp mass, with values computed per template bin and expressed in units of the standard deviation of the whitened strain data. This adaptive approach balances glitch suppression with sensitivity to real signals across the mass range (Sachdev et al. 2019; Ewing et al. 2024).

Along with gating, accurate whitening of strain data and template waveforms required PSD estimation using Welch's method with geometric-mean and median averaging to characterize frequency-dependent noise (Messick et al. 2017). In the online configuration, PSDs were estimated using a 4 s fast Fourier transform (FFT) computed continuously on incoming data and used to whiten the strain in real time.

To account for longer-term noise variations that could impact the SVD, template waveforms were re-whitened weekly using updated PSDs derived from recent data. For template whitening, the analysis period was divided into continuous segments of up to 8 h, with each segment producing a PSD using a 8 s FFT. The final PSD was constructed by taking the median power in each frequency bin across segments (Messick et al. 2017). The SVD basis for each template bin were then recomputed weekly using the same PSD. Weekly cadence balanced the need to track gradual detector noise evolution with the high computational cost of recomputing SVDs (Ewing et al. 2024).

After whitening the strain data using the PSD, matched-filtering was performed in the time domain, using the SVD basis, starting at 15 Hz for the stellar-mass region (10 Hz for the IMBH region). The output, an SNR time series, was scanned for peaks exceeding a threshold of 4, which were defined as triggers. A signal-consistency statistic, ξ^2, was computed for each trigger to quantify the deviation between observed and expected SNR values across a fixed number of points centered on the peak (Messick et al. 2017). Only the highest SNR trigger in each 1 s window per template bin was kept. Templates with chirp masses up to $1.73\,M_\odot$ used the TAYLORF2 approximant and 701 samples to compute ξ^2. Higher-mass templates used the SEOBNRV4_ROM approximant and 351 samples to compute ξ^2. The resulting triggers were used to form candidates for further processing.

Candidates were classified as either single-detector or coincident. Coincident candidates are defined as triggers from the same template in two or more detectors and occurring within the light-travel time between detectors, plus a 5 ms window to account for timing uncertainties. Single-detector candidates had no corresponding trigger in another detector within the window (Sachdev et al. 2019). Candidates from the IMBH region were excluded if they were a single-detector candidate. Each candidate was ranked using a LR quantifying the probability of astrophysical versus noise origin. The LR incorporated per-detector SNRs, the signal-consistency statistic ξ^2, local trigger rates, detector sensitivities, and prior assumptions about the CBC population. For coincident candidates, the relative arrival times and phases at the different detectors were included, with probabilities computed assuming isotropically-distributed sources. In contrast, single-detector candidates were down-weighted using a *singles-penalty*, empirically tuned via signal simulation campaigns, in order to suppress false positives while retaining sensitivity (Sachdev et al. 2019).

The LR was computed independently for each template bin to account for variations in how templates interact with detector noise. Background distributions for each template bin were built from single-detector triggers occurring during coincident observing time. Candidates were clustered in an 8 s window, and the candidate with the highest LR across all template bins was selected for further processing. The construction of the ranking statistic, including the modeling of the signal and noise distributions used in the LR calculation, is described in Cannon et al. (2015); Messick et al. (2017); Tsukada et al. (2023).

Candidate significance was quantified using FAR, defined as the rate at which detector noise could produce a candidate with equal or higher LR than the candidate under consideration. FAR was estimated by populating the background LR distribution via Monte Carlo sampling of SNR and ξ^2 values drawn from the empirical background, smoothed via kernel density estimation (KDE). Samples were assigned coalescence times, phases, and templates from uniform distributions (Cannon et al. 2013; Joshi et al. 2025a). To reduce contamination of the background by astrophysical signals, a 10 s window around each GW candidate is vetted and excluded

from the background, but only for candidates identified online by the online configuration (Joshi et al. 2023).

Dropped data from the online configuration was processed in an offline configuration, yielding additional candidates. Results from both configurations were merged, and FARs recomputed using the full background collected by the online configuration (Joshi et al. 2025b). The combined set of candidates and background defined the stellar-mass search analysis. To produce the final GWTC-5.0 results, candidates and background from the stellar-mass and IMBH search-analysis regions were combined into a unified analysis. Each region was assigned a weight, determined from injection campaigns (Essick et al. 2025), reflecting its relative sensitivity and the expected number of detectable astrophysical signals. The stellar-mass and IMBH search analyses were assigned weights of 0.94 and 0.06, respectively. The weights account for the fact that different analyses contribute unequally to the final set of candidates under the assumption of a given prior source population. FARs were rescaled by the inverse of these weights, and LRs were recomputed to provide a unified, unbiased ranking (Joshi et al. 2025a).

In addition to FAR, the total probability of astrophysical origin $p_{\rm astro}$, and the classification probability for the mutually-exclusive BNS, NSBH, or BBH source classes, were calculated for each candidate. Unlike other pipelines, GstLAL adopts a signal model that assumes the Salpeter primary mass function (Salpeter 1955) for each source class, with uniform distributions in mass ratio and spin. The source-specific probabilities were computed using a Bayesian framework that models the posterior probability distribution of the rates associated with each class. The framework treats GW triggers as realizations of independent Poisson processes, with rate estimates derived from previous observing runs and the population properties of each source class (Ray et al. 2023).

The GstLAL offline analysis of O4b included corrections from O4a which affected the ranking statistic and the calculation of $p_{\rm astro}$. The first change revised the handling of the horizon distance for each operating detector which is dependent on the detector configuration while observing (Tsukada et al. 2023). This corrected our unintended loss of sensitivity in the construction of the horizon distance. The second correction excluded the online vetted GW events from the background estimation during the re-ranking steps (Joshi et al. 2025b). This prevents the contamination of the background and bias FAR assignments (Joshi et al. 2023). Finally, the normalization used to map ranking information to $p_{\rm astro}$ was revised. This revision increases the odds-ratio, $p_{\rm astro}/(1-p_{\rm astro})$, by a factor of ~ 3.6, indicating that the previous normalization systematically underestimated $p_{\rm astro}$ values. The GstLAL search configuration, template bank design, and candidate ranking used for O4a and O4b differ from those used for O3. While GWTC-4.0 and GWTC-5.0 results were produced by using online results supplemented with offline analyses to avoid redundant processing, the GWTC-3.0 results were produced entirely from a post-run offline analysis (Abbott et al. 2023a). The O3 template bank extended to detector-frame total masses of $758\,M_\odot$, with aligned-spin parameters in $[-0.999, 0.999]$ for components more massive than $3\,M_\odot$, and was constructed using a stochastic placement algorithm (Abbott et al. 2021a). No duration constraint was imposed during matched filtering with the O3 template bank. Ranking of candidates identified by matched-filtering used an LR similar to O4a and O4b, but also incorporated the iDQ glitch likelihood (Godwin et al. 2020; Abbott et al. 2021a). The iDQ glitch likelihood was used only for single-detector candidates prior to GWTC-2.1, and the singles-penalty parameter, used to down-weight such triggers, was tuned differently than in O4a and O4b. Additionally, the LR term modeling the distribution of ξ^2 used an empirically-tuned analytical function in O3, rather than being directly informed by the statistical properties of ξ^2 as done in O4a and O4b. Furthermore, data near vetted GW candidates were not excluded from background estimation before O4a.

The GstLAL pipeline used during O2 differed from O3 primarily in the template bank construction and candidate ranking methodology. The O2 template bank covered detector-frame total masses up to $400\,M_\odot$ and mass ratios from 1/98 to 1, with aligned-spin parameters in $[-0.999, 0.999]$ for components above $2\,M_\odot$ (Abbott et al. 2019a; Mukherjee et al. 2021). The LR in O2 assumed uniform signal recovery across the template bank, ignoring the non-uniform template density and any astrophysical prior on the source population (Abbott et al. 2021a; Fong 2018).

The O2 GstLAL pipeline differed from O1 in its template bank construction, candidate ranking methodology and data conditioning. The O1 template bank covered detector-frame total masses from $2\,M_\odot$ to $100\,M_\odot$ (Mukherjee et al. 2021). In O1, the LR was computed only for coincident candidates and excluded time and phase differences between detectors (Hanna et al. 2020). The LR included the joint probability of observed SNRs but assumed equal horizon distances across detectors: an approximation that neglected how detector-specific noise characteristics impact the horizon distance. It was improved in O2 by pre-computing joint SNR distributions for a discrete set of horizon-distance ratios. Additionally, the O1 pipeline also lacked a template-mass-dependent glitch excision threshold (Sachdev et al. 2019).

3.3. MBTA

MBTA (Abadie et al. 2012; Adams et al. 2016; Aubin et al. 2021; Alléné et al. 2025) has performed online searches for CBCs since the initial-detector era science runs (Abadie et al. 2012). MBTA analyzes each detector independently using matched-filtering, before searching for candidates seen in coincidence. To reduce the computational cost, the pipeline splits the matched-filtering process over several frequency bands (typically two). Since O3, MBTA runs offline search analyses as well, with the aim of more accurately assessing the significance of candidates and contributing to the GWTC (Abbott et al. 2024, 2023a).

The parameter space explored by MBTA's template bank is entirely defined by the detector-frame masses and spins (assumed parallel to the orbital angular momentum) of the

binary components. The O4b analysis covers individual masses greater than $1\,M_\odot$, total masses up to $500\,M_\odot$, and mass ratios ranging from $1/20$ to 1.

In line with astrophysical expectations for galactic NSs (Stovall et al. 2018), spin magnitudes are restricted to 0.05 for objects with masses below $2\,M_\odot$, while more massive objects are allowed to have spin magnitudes of up to 0.998.

Very short templates, which tend to generate a background excess, are removed by introducing a cut-off of 200 ms on the template duration, calculated starting from 18 Hz. A visual representation of the parameter space covered by the O4b MBTA stellar-mass bank is given in the upper-right part of Figure 2.

The template banks used for O4 are generated with different algorithms. The stellar-mass search space is first divided into two regions, corresponding to BNS and symmetrical BBH (no more than a factor 3 between the two component masses). Templates in the BNS region are geometrically placed according to a local metric estimate (Brown et al. 2013), aiming for SNR recovery of at least 98%. Symmetrical BBH template placement uses a hybrid algorithm (Roy et al. 2017, 2019), with a similar expected SNR recovery. The rest of the bank is then completed using the same algorithm, with the objective of limiting the maximum SNR loss to 3.5%. As MBTA filters data in two disjoint frequency bands, a bank has been created for each of them, as well as one for the whole frequency range. Each template from the latter bank is associated with a template from other banks, maximizing their overlap in their common frequency band.

MBTA uses a total of about 942000 templates. The shortest 3000 templates are processed in the full-frequency band due to their intrinsically narrow bandwidth. For the longer templates, the two-band decomposition leads to approximately 50000 low-frequency-band templates and 20000 high-frequency-band templates. Filtering with such a decomposition leads to an SNR loss that is much smaller than the minimal match of the bank.

Templates use the SPINTAYLORT4 approximant (Klein et al. 2014) for total masses below $4\,M_\odot$ and SEOBNRV4_OPT (Bohé et al. 2017) above.

Before matched filtering, MBTA applies multiple preprocessing steps to the calibrated strain data delivered by the detectors. Data are first resampled to 4096 Hz. Next, a gating procedure is applied, designed to remove short and high-amplitude glitches. The gating used for O4 is similar to the version deployed for O3 and is triggered by drops in the detector sensitivity (Aubin et al. 2021). It is also applied to periods of poor quality data identified by data-quality vetoes (Abbott et al. 2020a). Less than 0.1% of the data provided by each detector was removed by MBTA gating during O4b (Alléné et al. 2025). An ungated search analysis is also performed for part of the parameter space, with the aim of finding the massive, intense CBCs that are likely to trigger the gating (Aubin et al. 2021).

MBTA regularly re-estimates the noise PSD of each detector by taking the median over thousands of seconds of data from several FFTs, whose lengths range from seconds to hundreds of seconds, depending on the targeted region of the parameter space (Aubin et al. 2025).

For most templates, the matched-filter calculation is carried out in two disjoint frequency bands, starting at 24 Hz and separated at 80 Hz. When a sufficiently high SNR is detected, bands are then coherently recombined to produce a unified single-detector trigger from the two bands. The shortest templates, however, are processed using a single band starting at 20 Hz due to their intrinsically narrower frequency band.

Triggers with an SNR above 4.4 (or 9 for the ungated analysis) are recorded. MBTA then applies several tools to reject triggers caused by transient noise. A basic χ^2 test discards any trigger whose SNR is not distributed as expected across the filtered bands (Adams et al. 2016). A second χ^2 statistic is used to quantify the discrepancy between the measured SNR time series and the template autocorrelation, and to reweight the SNR, thereby defining a single-detector ranking statistic (Aubin et al. 2021). MBTA also downgrades the ranking statistics of triggers occurring during periods when detectors show signs of poor data quality. It builds a χ^2 statistic from the differences between the time series of the maximum SNR observed around each trigger and models based on a large population of simulated signals (Alléné et al. 2025).

With the addition of Virgo in O4b (absent during O4a), MBTA now looks for coincident triggers between the two LIGO detectors and Virgo detector sharing the same templates, within a limited time window. A combined ranking statistic is computed as the quadrature sum of single detector ranking statistics. This statistic also includes a term measuring the consistency of arrival times and phases across the detectors (Aubin et al. 2021). To mitigate the impact of loud glitches on a single detector, paired with random fluctuations from other detectors, the pipeline rejects high-mass coincidences resulting from a heavily reweighted trigger and a low SNR trigger. In addition to the search for coincident triggers between the two LIGO and Virgo detectors, since the start of O4, MBTA now also produces single-Hanford and single-Livingston candidates. In O4b, Virgo is only used for coincidences. In order to limit the risk of false alarms, while maximizing the chances of detecting electromagnetically-bright sources, the MBTA search for single-Hanford and single-Livingston candidates is restricted to chirp masses below $7\,M_\odot$.

The significance of the candidates obtained by the stellar-mass search analysis is evaluated based on their probable origin, quantified by $p_{\rm astro}$. These probabilities are obtained using a method similar to the O3 method (Andres et al. 2022). For a given astrophysical source, the ratio of expected number of foreground triggers to the total number of triggers (foreground + background) is calculated. The foreground is estimated using population models similar to those used to estimate search sensitivity (see Section 3.7) and by fitting the rate observed by MBTA. The background is obtained by fitting the distribution of the ranking statistic measured by MBTA, after eliminating confidently-detected astrophysical signals. In the case of coincidences, the statistic is increased by considering fake coincidences between noise trig-

gers, regardless of the time of arrival. These operations are performed over several bins in the mass and spin parameter space covered by the search.

The pipeline associates a FAR to candidates using a method similar to the $p_{\rm astro}$ background estimation. In order to make the values of $p_{\rm astro}$ and FAR consistent with each other, and to include an astrophysical prior in the FAR definition, $p_{\rm astro}$ is used since O4a as a new ranking statistic.

MBTA performed several online search analyses during O4b using a pipeline version similar to the O4a offline analysis (Abac et al. 2025b), in addition to searching for coincidences with Virgo data. All coincident candidates, as well as single-Hanford and single-Livingston candidates with chirp mass below $7\,M_\odot$, from the full-bandwidth analysis with a FARs below $2\,\mathrm{h}^{-1}$ were submitted to GRACEDB with a median latency of 21 s. The online $p_{\rm astro}$ calculation performed by the pipeline during O4b used the same model as for the offline O4a analysis (Abbott et al. 2023a), adjusting the signal rates to account for the increased sensitivity. MBTA was also conducting *early-warning* searches for CBCs with component masses between $1\,M_\odot$ and $2.5\,M_\odot$ and aligned spins below 0.05, with the aim of rapidly identifying high-SNR signals.

MBTA also conducted offline search analyses during the O3 and O4a period, and thus participated in the production of the previous catalog versions GWTC-2.1 (Abbott et al. 2024), GWTC-3.0 (Abbott et al. 2023a) and GWTC-4.0 (Abac et al. 2025c). During O4a, MBTA was operated in a similar configuration than for O4b, with the main difference of searching for mass ratio down to $1/50$. During O3, the pipeline processing was slightly different. The O3 stellar-mass search template bank was created using a purely stochastic method (Privitera et al. 2014), guaranteeing a SNR loss below 3%, compared with 2% for parts of the O4 banks. During O3, the pipeline only searched for CBCs with (detector-frame) total mass below $200\,M_\odot$, but no constraints on waveform duration were applied. Re-ranking of triggers occurring during periods of poor data quality was handled by comparing the trigger rate before and after the application of various other rejection criteria. Before O4, MBTA only calculated significance for candidates seen in at least two detectors. The FAR values were based on a ranking that assumed all templates were equiprobable, which was responsible for some inconsistencies with the associated $p_{\rm astro}$ values. Trials factors were used to account for the various coincidence types and parameter space regions (Aubin et al. 2021).

3.4. PYCBC

The PYCBC search pipeline (Dal Canton et al. 2014; Usman et al. 2016; Nitz et al. 2017) is a descendant of the IHOPE pipeline (Allen et al. 2012; Babak et al. 2013), which was used to search LIGO–Virgo data during the initial detector era up to 2010 (Abadie et al. 2012; Aasi et al. 2013). The input to PYCBC is the calibrated strain data from the detectors, which undergo a series of conditioning steps to prepare for analysis. After removing invalid or contaminated strain data flagged by CAT1 vetoes, we then apply high-pass filtering to suppress low-frequency noise (below 15 Hz), downsampling from 16 kHz to 2 kHz to reduce data volume, and gating (windowing out) to remove loud glitches (Usman et al. 2016). The pipeline is also limited to analyze data segments with a minimum length of 500 s, to ensure a sufficiently precise PSD estimate.

After data selection and input conditioning, data from each detector are filtered using a fixed template bank. The templates are placed in a four-dimensional parameter space of component detector-frame masses and orbit-aligned spins. For O4a, the template bank was generated with a hybrid geometric-random placement algorithm (Roy et al. 2017), imposing a minimum template waveform duration of 70 ms which allows coverage of systems with IMBHs, and applying a dynamic minimal match criterion to ensure smooth template density across the mass range (Roy et al. 2017, 2019). The O4b analysis reused the template bank generated during O4a. The match was calculated assuming the projected LIGO O4 sensitivity that predicts a 160 Mpc BNS range (LIGO Scientific Collaboration et al. 2022). The waveform model used is SEOBNRV5_ROM (Pompili et al. 2023) for total mass above $4\,M_\odot$, and TAYLORF2 (Vines et al. 2011) below. A fixed lower cutoff frequency of 15 Hz is applied for systems with total mass above $100\,M_\odot$. The bank covers a total mass range of $2\,M_\odot$ to $500\,M_\odot$, with mass ratios from $1/97.989$ to 1, and aligned spin ranges from -0.997 to 0.997 for BH components and from -0.05 to 0.05 for NS components, as shown in Figure 2.

Data from each detector are independently matched filtered against the template bank, producing an SNR time series for each template: triggers, corresponding to potential signals, are then identified as local SNR maxima above a given threshold within predefined time windows. A χ^2 test is applied to assess whether the time–frequency distribution of power in the data matches the expected distribution from the template waveform (Allen 2005). The pipeline re-weights the matched-filter SNR using the reduced χ^2 normalized such that the expected value in Gaussian noise for a signal matching the template is unity (Abadie et al. 2012; Usman et al. 2016). Triggers with re-weighted SNR below a threshold of 4 are discarded as they are likely to correspond to noise transients. Since short-duration blip glitches (Cabero et al. 2019) may pass the time–frequency χ^2 test for some templates, a high-frequency sine–Gaussian χ^2 test is used to further identify such glitches (Nitz 2018).

A single-detector ranking statistic $\hat{\rho}$ is then calculated, incorporating the time–frequency and sine–Gaussian χ^2 tests, as well as a correction to the SNR due to short-term PSD variability (Zackay et al. 2021; Mozzon et al. 2020).

The next step of the pipeline is to identify candidates consistent with potential GW signals by comparing the coalescence times and template parameters of triggers from multiple detectors. Coincident candidates are found if triggers from two or more detectors, associated with the same template, occur within a time window that accounts for the light travel time between detectors and a small margin for timing errors. If data from more than two detectors are analyzed,

coincident candidates may be formed across multiple combinations of two and three detectors via the same process.

Starting in O4, candidates consisting of triggers in only a single detector are also included (Cabourn Davies & Harry 2022), though only from templates with duration larger than 0.3 s and satisfying additional cuts related to the background distribution.

Following this, the pipeline calculates the statistical significance of candidates by estimating their FAR. A ranking statistic is assigned to each candidate, reflecting its likelihood of being a true GW signal versus background noise. For coincident candidates seen in two or three detectors, the FAR is computed by comparing their ranking statistics to a set of artificially generated background candidates created by time-shifting triggers in one of the participating detectors. In O4a, where only LHO and LLO data were available, background candidates were generated by repeatedly shifting triggers in one detector with respect to the other by a fixed interval. In O4b, incorporating data from Virgo, background candidates are generated by extending the time-shifting procedure to a multi-detector scheme (Davies et al. 2020). Artificial triple-detector background coincidences are generated by identifying coincident triggers in any two detectors and shifting the triggers in the third detector by fixed intervals. For three-detector analyses, the FAR is computed by summing the FAR contributions from various coincidence types at a given ranking statistic threshold: triple-detector (HLV) and the three possible double-detector combinations (HL, HV, and LV) (Davies et al. 2020).

To ensure that candidates are statistically independent, clustering is performed on both coincident (un-shifted) and time-shifted analyses: only the candidate with the highest detection value within a set time window is kept. To reduce background contamination due to loud signals, any candidate detected with FAR below a threshold of 1 per 100 yr is removed from the background estimation for less significant candidates (Abbott et al. 2016b; Nitz et al. 2019).

For single-detector candidates in O4, since their FARs cannot be estimated from time-shifted analysis, the ranking statistic distribution is fitted with a falling exponential, and extrapolated with a maximum possible inverse false-alarm rate (IFAR) assignment of 1000 yr (Cabourn Davies & Harry 2022). In times when multiple detectors are active, FAR estimates from all candidate types at the ranking statistic value of the candidate are added to give the final FAR estimate.

The ranking statistic assigned to both coincident and single-detector candidates is designed to optimize the detection rate by reflecting the relative densities of signal vs. noise over the binary source parameter space. To account for noise variations across different templates, the single-detector background distributions are fitted, with the fit parameters allowed to vary over the template parameter space (Nitz et al. 2017). We define an intermediate expression R_3, which is the basis for both the multi-detector and single-detector statistic:

$$\begin{aligned} R_3 &= \ln\left(\frac{\sigma^3_{\mathrm{min},i}/\bar{\sigma}^3_{\mathrm{HL},i}}{A_{\mathrm{N}\{d\}}\prod_d r_{d,i}(\hat{\rho}_d)}\frac{p(\mathbf{\Omega}|\mathrm{S})}{p(\mathbf{\Omega}|\mathrm{N})}\right) \\ &= \ln p(\mathbf{\Omega}|\mathrm{S}) - \ln p(\mathbf{\Omega}|\mathrm{N}) - \ln A_{\mathrm{N}\{d\}} \\ &\quad - \sum_d \ln r_{d,i}(\hat{\rho}_d) + 3(\ln\sigma_{\mathrm{min},i} - \ln\bar{\sigma}_{\mathrm{HL},i}). \end{aligned} \tag{6}$$

Here $p(\mathbf{\Omega}|\mathrm{S})$ and $p(\mathbf{\Omega}|\mathrm{N})$ are the probabilities of a signal or noise candidate, respectively, to have extrinsic parameters $\mathbf{\Omega}$, comprising relative phases, arrival times and amplitudes between detectors (Nitz et al. 2017). We consider the noise distribution $p(\mathbf{\Omega}|\mathrm{N})$ as uniform for a given combination of detectors, absorbing this constant factor into the normalization of $p(\mathbf{\Omega}|\mathrm{S})$. The rate of noise events is predicted from the allowed time window $A_{\mathrm{N}\{d\}}$ for coincident events involving detectors $\{d\}$ and the expected rate of noise triggers $r_{d,i}(\hat{\rho}_d)$ in template i at re-weighted SNR $\hat{\rho}_d$. The last term accounts for the time-dependent rate of signals, which is proportional to the network sensitive volume for a given template, and scales with the cube of the minimum sensitive distance $\sigma_{\mathrm{min},i}$ over triggering detectors, normalized to an average LIGO sensitive distance $\bar{\sigma}_{\mathrm{HL},i}$.

The ranking statistic used in O4 includes additional terms beyond Equation (6). First, an explicit model of the signal distribution over binary masses and spins is incorporated, intended to optimize the detection rate of the known CBC population while remaining sensitive to so far unpopulated parameter regions (Kumar & Dent 2024). Second, the noise event rate model is updated to account for variations in the rate of noise artifacts, correlated with auxiliary channels that monitor the detector state or environmental disturbances, as measured by IDQ (Essick et al. 2020). We then have:

$$\begin{aligned} R_{\mathrm{O4}} &= R_3 + \ln d_{\mathrm{S}}(\boldsymbol{\theta}_{\mathrm{int}}) - \ln d_{\mathrm{N}}(\boldsymbol{\theta}_{\mathrm{int}}) \\ &\quad - \sum_d \ln\delta_d\left(\Delta_d(t), B_d(\boldsymbol{\theta}_{\mathrm{int}})\right) \\ &\quad + \sum_D \ln C_D, \end{aligned} \tag{7}$$

where $d_{\mathrm{S}}(\boldsymbol{\theta}_{\mathrm{int}})$ and $d_{\mathrm{N}}(\boldsymbol{\theta}_{\mathrm{int}})$ represent a KDE-based model distribution of signals over template parameter space $\boldsymbol{\theta}_{\mathrm{int}}$, and the distribution of template points, respectively. The time-dependent excess rate of non-Gaussian triggers is notated as $\delta_d(\Delta_d(t), B_d(\boldsymbol{\theta}_{\mathrm{int}})) \geq 1$ for a detector d. Our estimate of δ_d depends on a discrete detector state variable $\Delta_d(t)$ and on an index $B_d(\boldsymbol{\theta}_{\mathrm{int}})$ labelling template bins. We consider three possible detector states, corresponding to (i) auxiliary channels indicating significant likelihood of noise artefacts; (ii) presence of loud gated glitches nearby in time, or (iii) neither. In the case of single-detector candidates, the ranking statistic omits terms that depend on multi-detector coincidence quantities (Cabourn Davies & Harry 2022); specifically, the terms $A_{\mathrm{N}\{d\}}$, $p(\mathbf{\Omega}|\mathrm{S})$, and $p(\mathbf{\Omega}|\mathrm{N})$ present in Equation (6) are omitted.

The remaining term C_D is a constant correction dependent on the candidate type D (a single detector or combination of detectors). It accounts for the different rates of triggers due to detected signals and noise over the available combinations of detectors, after other terms in Eq. (7) are applied, with the aim of improving the overall search sensitivity. For the GWTC-4.0 analysis of O4a data Abac et al. (2025b), this term was not included. For the analysis of O4b data and a revised O4a analysis, both included in GWTC-5.0 (Abac et al. 2026b), corrections are applied to all candidates, apart from LHO–LLO double coincidences (i.e., we impose $C_{\rm HL} = 0$). The values of C_D are empirically tuned using simulated signals (injections) over representative periods of O4 data. Inclusion of the correction for O4a reanalysis resulted in an increase of up to 10% in PyCBC's sensitivity (at a threshold of FAR $< 1\mathrm{yr}^{-1}$).

The probability of astrophysical origin $p_{\rm astro}$ is estimated by a Poisson mixture inference (Abbott et al. 2016a,c), for which the distribution of signal events over the ranking statistic is estimated via histograms from a set of simulated signals and distributions of noise events are estimated via histograms of time-shifted background for coincident candidates, or via exponential fits for single-detector candidates (Cabourn Davies & Harry 2022). Such histograms are generated separately for each type of (single-detector or coincident) candidate in single-detector or multi-detector observing time to allow for differences in event rates and distributions (Dent 2025). Marginalization over the unknown rate of astrophysical signals is performed by generalized Gauss–Laguerre quadrature (Creighton 2019). The resulting candidate $p_{\rm astro}$ values are apportioned between different astrophysical categories using an estimate of the source chirp mass based on search outputs (Dal Canton et al. 2021; Villa-Ortega et al. 2022).

Several features of the PyCBC online search analysis (Dal Canton et al. 2021; Nitz et al. 2018) differ from the offline search analysis, these are detailed in Section 3.4 of Abac et al. (2025b). The main differences between O4a and O4b PyCBC online analysis, are:

- Improved gating: to improve glitch mitigation, the gating window was extended to a data segment significantly longer than the analysis chunk (16 s). This ensures that extended clusters of multiple glitches are detected and gated across multiple chunks, preventing them from persisting through the duration of the longest signal templates.

- Virgo data was used in the end-to-end offline analysis. For online analysis Virgo data was not used to perform low-latency matched-filtering but only for localization when at least one other detector in the network registered a trigger.

- Deploying an improved version of the SNR optimizer used for O3 analysis (Abbott et al. 2023a): for every trigger a refined search near to the triggering template is performed to recover the potentially lost SNR due to the discreteness of the template bank. The O4b SNR optimizer is tuned to improve upon the settings (in particular the range of allowed masses) used in the previous O3 analysis.

The above descriptions concern PyCBC analyses of O4a and O4b data. Here, for completeness, we summarize significant changes in search methods used for archival catalog (offline) results since the GWTC-1.0 release. The O3 template bank was constructed using the same hybrid geometric-random method as in O4 to ensure efficient coverage of the parameter space (Roy et al. 2017). Three different independent search analyses (PyCBC-BBH, PyCBC-Broad, and PyCBC-IMBH) were used, with each covering distinct regions of the parameter space (Abbott et al. 2023a; Chandra et al. 2021). The waveform model SEOBNRv4_ROM was used for systems above $4\,M_\odot$, and TaylorF2 for lower-mass systems (Roy et al. 2019). Templates with durations below $0.15\,\mathrm{s}$ were excluded to reduce false alarms from short glitches, thereby also placing an upper limit on the mass of the systems, unlike in O4, where shorter-duration templates are included to improve coverage of high-mass systems. The template bank used in the earlier O2 and O1 analyses (Dal Canton & Harry 2017) spanned total masses from $2\,M_\odot$ to $500\,M_\odot$, with mass ratios from $1/98$ to 1, and spin limits based on component mass: 0.05 for masses below $2\,M_\odot$ and up to 0.998 for higher masses (Abbott et al. 2019a). For O3, the background estimation was extended to a three-detector network by using fixed relative time shifts between detectors (Davies et al. 2020), whereas in earlier runs, a two-detector method (also used in O4a) was employed. Unlike in O4, single-detector candidates were not included in O3 or previous runs. The O3 broad search statistic is given by Equation (6), i.e., the O4 statistic without the additional KDE and data-quality terms. For the O3-BBH search, an alternate ranking statistic was adopted, incorporating an additional term in the broad search statistic dependent on the template chirp mass to enhance sensitivity to the observed BBH population (Nitz et al. 2020). In earlier runs, only one type of event (LHO–LLO) was considered, and sensitivity variations across the parameter space, such as those described by σ_i, were neglected. As a result, many terms in the O3 statistic in Equation (6) either disappeared or became constants and were omitted (Nitz et al. 2017).

3.5. SPIIR

SPIIR is a coherent search pipeline developed primarily for the online detection of GW signals from CBCs (Hooper et al. 2012; Luan et al. 2012; Chu et al. 2022). It became operational during O3, during which it registered 38 out of the 56 non-retracted public alerts (Abbott et al. 2021a, 2024; Chu et al. 2022). The design of SPIIR focuses on efficiency and rapid processing, enabling it to provide detections with latencies around 11 s. This capability makes SPIIR particularly useful for issuing early warning alerts for electromagnetic follow-ups of GW candidates (Kovalam et al. 2022). While an offline version of SPIIR is under development, it

has not yet participated in any of the offline searches through the end of O4b.

The SPIIR pipeline employs infinite impulse-response (IIR) filters as template banks to perform matched filtering directly in the time domain (Hooper et al. 2012). Unlike traditional methods, which can be time-consuming, the IIR filtering approach allows SPIIR to execute matched filtering rapidly. The primary advantage of IIR filters is their ability to approximate matched-filtering SNR in a way that can be efficiently parallelized and executed on graphics processing units (GPUs), which significantly accelerate high-performance computing tasks (Liu et al. 2012; Guo et al. 2018). The SPIIR template-bank generation is a two-step process. The first step is to generate the template points in the CBC parameter space using the stochastic placement algorithm (Privitera et al. 2014; Harry et al. 2009) with a minimum match of 0.97. The approximant SpinTaylorT4 (Klein et al. 2014) is used for chirp mass below $1.73\,M_\odot$ and SEOBNRv4_ROM for larger masses. In the second step, each corresponding waveform is decomposed into approximately 350 IIR filters (Hooper et al. 2012; Liu et al. 2012; Guo et al. 2018). These IIR filters are then utilized for matched filtering using GPU parallelization.

The real-time data analysis process begins by retrieving strain data for all detectors in observing mode.

The data are then conditioned, applying data-quality vetoes, down-sampling from 16 384 Hz to 2048 Hz, and whitening, which requires real-time PSD calculation. The PSD calculation takes place using the Welch method with tens of 4 s overlapping data blocks. This method prevents sporadic transient signals from biasing the PSD estimate. The data are also subjected to gating to mitigate the effects of loud transient noises. Once conditioned, the whitened data from each detector is passed through a set of IIR filters that span the parameter space of CBC templates, generating corresponding SNR time series for each template.

A key distinguishing feature of the SPIIR pipeline is its coherent search algorithm (Bose et al. 2011; Harry & Fairhurst 2011). The coherent search algorithm looks for time and phase consistency across all detectors. SPIIR calculates a maximum network likelihood ratio statistic to evaluate the likelihood of a coherent GW signal being present in the data. This approach enhances the pipeline's sensitivity, particularly to weaker signals, and improves localization accuracy by combining data from multiple detectors.

The coherent analysis can be computationally demanding. To address this challenge, SPIIR uses SVD (Wen 2008), which allows for efficient calculation of the coherent statistic by focusing on the most significant signal components. This optimization reduces the complexity of the search and enables rapid scanning across different sky positions to localize the source.

After determining the coherent SNR, SPIIR applies a ξ^2 signal-consistency test (Messick et al. 2017) to asses the morphological consistency between data and template and combines it with the coherent SNR into a multidimensional ranking statistic. FARs are assigned by comparing candidates with background distributions generated from time-shifted analyses. To ensure robust estimation, the pipeline collects at least one million background candidates, typically accumulated over several hours.

FARs are assigned using a minimum of one week of background data. The background is extrapolated using a k-nearest-neighbor (KNN) KDE method with $k = 11$, smoothing the probability density using neighboring bins. This approach allows accurate FAR estimation even for rare, high-significance candidates. To account for non-stationary noise, SPIIR additionally maintains shorter-term background sets over 2 h and 1 d. Furthermore, it stores individual single-detector SNR and ξ^2 values and applies the same KNN-based extrapolation method to estimate single-detector FARs used in veto stages before making the final decision about the trigger's validity.

In preparation for O4, SPIIR has undergone several upgrades, with some already implemented and others still in testing and review. These improvements include: automated signal trigger removal from the background estimation, preventing astrophysical triggers from affecting the background calculation; fine tuning of the signal consistency test ξ^2, aiming to improve sensitivity to BNS and NSBH; and a dynamic approach to background collection. Furthermore, SPIIR has implemented the necessary infrastructure to process KAGRA data.

3.6. *Criteria for Inclusion in the Catalog*

Search pipelines can in principle produce a rate of candidates as high as one every few seconds, depending on their configuration choices. With the sensitivity of current detectors, however, the vast majority of such candidates would have low significance and would be unlikely to have an astrophysical origin. For practical consideration, we therefore select a smaller number of candidates to report in the GWTC.

The criterion for selecting which candidates to include has evolved with the observing runs and GWTC versions, due to the rapid evolution of the field and the computational and human cost required to retroactively apply the latest analysis methods to previous observing runs. In GWTC-1.0 (Abbott et al. 2019a), we published all candidates from O1 and O2 for which at least one matched-filtering pipeline (PyCBC or GstLAL at the time) estimated a FAR below 1 per 30 d.

GWTC-2.1 (Abbott et al. 2024) and GWTC-3.0 (Abbott et al. 2023a) respectively added candidates from the first half of the third observing run (O3a) and second half of the third observing run (O3b) with a FAR below $2\,\mathrm{d}^{-1}$ in any pipeline, with the exception of SPIIR. Following this logic, candidates from O4a whose FAR is less than $2\,\mathrm{d}^{-1}$ in any of the pipelines described in Section 3, with the exception of SPIIR, were added to GWTC-4.0, and we now do the same for O4b and GWTC-5.0. In general, we define a combined candidate list (denoted as the *Any* pipeline results) for a given FAR threshold by including candidates whose FAR is below that threshold in at least one pipeline.

When candidates from multiple pipelines are reported with times within a given clustering window (LIGO–Virgo–

KAGRA Collaboration 2025b), then they are grouped under a single event name. FARs from individual pipelines are presented side-by-side and are neither corrected for multiple trials, nor combined in any way. Downstream analyses that need a single clustered list of candidates are then free to apply any appropriate combination method, e.g., choosing the minimum FAR across the available pipelines.

Stricter candidate selection criteria are typically applied when performing downstream analyses and for the purpose of displaying lists and properties of the candidates. Stricter criteria might, for example, be motivated by computational or person-power considerations. Furthermore, downstream analyses might not necessarily have the same requirements; for instance, a sufficiently large SNR might be necessary, as opposed to a sufficiently low FAR. Such considerations, and consequently the corresponding selection criteria, might in turn evolve as data-analysis methods become more efficient. Stricter criteria will therefore be described wherever necessary.

3.7. *Search Sensitivity*

Apart from the candidates themselves, an important product of search pipelines is a measurement of their sensitivity, which is necessary to diagnose the behavior of the pipelines and to infer the rate density of the astrophysical sources. We quantify the sensitivity of the search pipelines described above by their estimated time–volume product or hypervolume $\langle VT \rangle$ (Abbott et al. 2023a). The hypervolume represents the sensitivity of a given search to a set of sources assumed uniformly distributed in both comoving volume and source-frame time. An estimate of the total number of signals $\hat{N}$ that a pipeline is likely to detect during a period of observation is given by

$$\hat{N} = \langle VT \rangle \mathcal{R}, \tag{8}$$

with $\mathcal{R}$ the rate of CBCs per unit comoving volume and source-frame time (Abbott et al. 2016a). In practice, hypervolumes are obtained through weighted Monte Carlo simulations, which use simulated GW signals referred to as *injections*.

The injections are added to the strain data, which are then processed by the search pipelines to determine which injections are detectable according to a given set of selection criteria. The resulting $\langle VT \rangle$ estimate is a function of the selection criteria, which for GWTC-5.0 have evolved over the observing runs as described in Section 3.6. Downstream analyses that use the entirety of GWTC, such as rate density inferences, can and should account for this evolution by applying the appropriate time-dependent selection criteria when considering which injections are detectable.

Our current population of injected signals includes only CBCs, as no transient GW signal so far detected is inconsistent with a CBC source at high confidence. The parameter ranges and distribution models were chosen to represent CBC signals the detectors were possibly sensitive to, based on previous releases of GWTC (Abbott et al. 2023b). The rate of injections was chosen to be much higher than the astrophysical rate in order to reach a sufficient precision in $\langle VT \rangle$. As the production of injections required detector PSDs to be measured over each month of observation, injection analysis was performed offline. To ensure consistency in the $\langle VT \rangle$ measurements, all pipelines contributing GWTC candidates (GstLAL, MBTA, PyCBC, and cWB-BBH) analyzed the same sets of injections. The pipelines processed data with injections using methods equivalent to those used to produce candidates from injection-free data, described earlier in this section. All pipelines used background data collected by their injection-free analyses to assign significances to the recovered injections. The choice of population priors and the injection generation process are detailed in Essick et al. (2025), while results from the injection analysis are described in Abac et al. (2026b).

The sensitive $\langle VT \rangle$ differs between pipelines, reflecting their abilities to detect CBCs within a given range of parameters as well as their different effective live times. As cWB-BBH is designed to search for BBH signals as described in Section 3.1, for this pipeline we calculate $\langle VT \rangle$ only for masses corresponding to BBH sources.

There are differences in how certain pipelines process injections compared to injection-free data, which arise from computational efficiency or technical considerations and do not significantly affect the estimated sensitivity. The cWB-BBH pipeline performs injection analysis in a $\pm 5\,\mathrm{s}$ time window around the time of each injection. The GstLAL and PyCBC pipelines process data using subsets of their template banks based on the known chirp mass of the injections, in order to limit the computational cost. In addition, GstLAL's injection processing has two other major differences compared to its injection-free data processing. First, injections are processed entirely in an offline configuration, as opposed to reassigning significances to triggers obtained in online analysis, as injection data were not available during the online analysis. This is a difference from other pipelines, which process both injection-free and injection data in their offline configurations to produce the GWTC-5.0 results and $\langle VT \rangle$ estimations. Second, GstLAL computes SNR time series only over a short time window ($\sim 6\,\mathrm{s}$) around the time of the injections, to limit the computational cost, as opposed to computing the SNR time series for all available data in O4b.

The search sensitivities may be calculated at different significance thresholds, which enables us to self-consistently estimate the number of astrophysical signals among the subthreshold candidates. If the source population distribution assumed for p_{astro} calculations is close to the true astrophysical distribution, then for each individual pipeline the sum of p_{astro} values below the given threshold provides the expected number of true signals among the candidates produced by the pipeline. This expected signal count is itself a realization of a Poisson process with a mean proportional to the pipeline's $\langle VT \rangle$ for the true signal population. This reasoning also applies to the total set of candidates identified by at least one pipeline, i.e., to the catalog as a whole, if we consider the

Any pipeline sensitivity. Therefore, following a similar approach in GWTC-3.0 (Abac et al. 2025a), we scale the sum of $p_{\rm astro}$ values for each pipeline by the ratios of the *Any* pipeline $\langle VT \rangle$ to the individual pipeline's, each of which is measured by total number of recovered injections without any weighting in the mass distribution. While these ratios in principle depend on the FAR threshold at which we count injections as recovered, the dependence is small and we typically impose a threshold of $2\,\mathrm{d}^{-1}$. We here make use of the injected distribution (Essick et al. 2025) designed to broadly follow the astrophysical CBC distribution inferred from O3 observations described in Abbott et al. (2023b).

We thus obtain an estimate of the signal count over subthreshold candidates for the *Any* pipeline by scaling the sum of $p_{\rm astro}$ from each of the four pipelines; finally, we take the average of these estimates across the four pipelines. The purity of the subthreshold candidate set is this estimated signal count divided by the overall number of subthreshold candidates.

4. DATA-QUALITY VETTING AROUND CANDIDATES

A subset of the candidates produced by the search algorithms described in Section 3 undergo studies to understand the quality of the interferometric data surrounding the candidates with a goal to characterise and, if required, mitigate against any departures from stationary Gaussian noise such as glitches. This is vital to ensure that downstream analyses, particularly PE (see Figure 1 and Section 5), which typically assume stationary Gaussian noise, produce unbiased results (Pankow et al. 2018; Powell 2018; Ghonge et al. 2024). These methods are all applied after candidates are identified by search algorithms and to short sections of data. A discussion of the data-quality studies that feed into the search algorithms is given in Section 3.

Event validation is the process of identifying and examining potential data-quality issues in the vicinity of a GW candidate. Several distinct tests are used, due to the variety of data-quality issues that may impact the detection or analysis of GW candidates. For each GW candidate, test results are collected into a report referred to as a data-quality report (DQR). The creation of a single report simplifies the process of interpreting the results of multiple tests together and allows for summary statistics about individual GW candidates to be produced based on the results of multiple aggregated analyses. We apply this process to all candidates reported by online search pipelines and labeled as significant, with an internationally-coordinated team to provide rapid vetting aided by an automated DQR framework (Soni et al. 2025; Arnaud et al. 2026). Manual data-quality vetting of these online candidates was only conducted when at least one test identified a potential data-quality issue; if no issues were identified in this automated step, no additional vetting was conducted at this stage. The manual validation process is also completed for all candidates identified by offline search pipelines that meet the requirements for PE analyses.

In contrast to O3, where LIGO and Virgo conducted separate validation processes (Davis et al. 2022; Acernese et al. 2023a), O4 employs a unified infrastructure across all active detectors. This uniformity in validation improves efficiency, streamlines information flow between different analysis groups, and reduces the demand on human resources. The first step in event validation is the creation of a DQR, which includes numerous tests of strain data-quality as well as analyses of auxiliary sensor information. The auxiliary sensor data from each observatory is generally only available at the observatory site, requiring these DQR analyses to be conducted at multiple locations across the network. In O4b, there were four DQR nodes that rely on a combination of the DQRBUILD (Davis et al. 2026) toolkit and VIRGO DQR (Acernese et al. 2023b): one at each of the LIGO Hanford, LIGO Livingston, and Virgo observatories, and a central node that hosts a database tracking results from all nodes.

Each DQR node is launched in response to an IGWN-ALERT;[1] the corresponding data-quality analyses are managed by a DQRBUILD workflow configurator at the central node and at LIGO Hanford and LIGO Livingston, and by VIRGO DQR at Virgo. Once results are generated, they are then uploaded to the central node to be tracked by a DQRBUILD SQL database. This database is used to generate a centralized result page that tracks the global set of data-quality analyses conducted for each event to simplify the manual event validation process.

The data-quality analyses hosted by each DQR node include tests of the Gaussianity and stationarity of the strain data residual (Vazsonyi & Davis 2023; Mozzon et al. 2020; Di Renzo et al. 2024), machine-learning signal versus noise classification (Álvarez-López et al. 2024), estimates of the environmental coupling (Helmling-Cornell et al. 2024), statistical correlations with auxiliary sensors (Essick et al. 2020; Smith et al. 2011; Vajente 2022; Isogai et al. 2010), tests of anomalies in the long-term behavior of the detectors, checks of known data-quality issues, and investigations of the detector observing state. Analyses that are statistical in nature automatically produce conclusions about the data quality of a candidate; however, the results from many tasks still require manual investigation by data-quality experts.

Once the initial online vetting is completed, all significant candidates undergo further scrutiny to determine a final assessment of the data-quality status. In this stage, tracked with the EVENT-VALIDATION package (Soni et al. 2025; Di Renzo & Virgo Collaboration 2023), a data-quality expert reaches a conclusion based on the preponderance of evidence in the automated DQR analysis results. If data-quality issues are identified in this stage of the process, additional investigation and noise mitigation efforts are conducted to ensure that these issues have minimal impact on estimates of the astrophysical parameters of candidates.

If data-quality issues, such as the presence of glitches, are identified around a candidate and these issues are not severe enough to warrant a retraction, we quantify if further action is required by comparing the distribution of excess power in

[1] https://igwn-alert.readthedocs.io

the relevant region with expected Gaussian noise and computing a p-value (Vazsonyi & Davis 2023). Until November 2023, a conservative threshold of 0.1 was adopted. This threshold proved to result in an unsustainable rate of false positives, and was therefore changed to 0.05. For p-values greater than this threshold, we conclude the data are consistent with Gaussian noise, and the candidates are ready for downstream analyses. On the other hand, if $p \leq 0.05$, we attempt to subtract the noise using either a modeled Bayesian inference approach implemented in BAYESWAVE (Pankow et al. 2018; Cornish et al. 2021; Chatziioannou et al. 2021; Hourihane et al. 2022; Gupta & Cornish 2024), or a linear noise subtraction based on auxiliary witness channels (Davis et al. 2022). For all candidates in O4b, BAYESWAVE was used and applied to the `GDS-CALIB_STRAIN_CLEAN_AR` channel as described in Abac et al. (2026c); no correction is made for calibration uncertainty in this process. For past observing runs, BAYESWAVE has predominantly been applied, except in cases where a witness channel was available (Davis et al. 2022; Abbott et al. 2023a). However, if the noise is extended in time or frequency, as is often the case for certain glitch classes (e.g., Hourihane et al. 2022; Soni et al. 2021; Soni et al. 2024), making subtraction difficult, or if the data remains insufficiently stationary after subtraction, we instead restrict the time and frequency analysis window to avoid the afflicted region. For badly afflicted candidates, preliminary PE studies are routinely performed to validate the mitigation technique and understand the impact on parameter estimates (Davis et al. 2022). Once a satisfactory mitigation approach has been identified, the candidates are ready for downstream analyses. In Abac et al. (2026b), we provide a summary of new GWTC-5.0 candidates that required mitigation and details of the mitigation approach that was applied. We also release the glitch-mitigated strain data as part of the GWTC-5.0 data release (LIGO–Virgo–KAGRA Collaboration 2026).

5. PARAMETER ESTIMATION FOR CBCs

For a selection of candidates identified by the search algorithms (see Section 3) and either having passed validation or with appropriate mitigation (see Section 4), we use Bayesian inference to estimate the source parameters $\boldsymbol{\theta}$ of the GW signal (see Figure 1), which includes both the intrinsic parameters of the CBC source and the extrinsic parameters that localize and orient the source in spacetime (Abac et al. 2026a). These inferences come in the form of posterior distributions $p(\boldsymbol{\theta}|d)$ over the source parameters. The posteriors represent our best understanding of the properties of the source and its location, including all correlations and degeneracies in the measured parameters for a given set of prior assumptions. The posteriors for the selected GW candidates from O4b are reported for the first time in Abac et al. (2026b) and can be obtained from the GWTC-5.0 data release (LIGO–Virgo–KAGRA Collaboration 2025d). We have not updated our inferences for candidates identified in previous observing runs. Thus, our preferred parameter inferences for candidates identified before O4b remain the same as previously reported in Abbott et al. (2024, 2023a); Abac et al. (2025a), with the exception of the BNS candidate GW170817 whose preferred parameter inferences are reported in Abbott et al. (2019a).

In Sections 5.2 to 5.8, we describe the broad elements of GW PE and discuss the particular choices made in each version of the catalog in Section 5.9, noting that configuration files to reproduce the analysis of individual candidates are provided with the data release.

5.1. *Bayesian Formalism*

To carry out Bayesian PE, we assume the data d contains colored Gaussian noise and an astrophysical GW signal which is well-approximated by a quasi-circular CBC waveform model (see Section 2), consistent with GR and absent of environmental effects. These assumptions determine the form of our likelihood $p(d|\boldsymbol{\theta})$, which for a single detector is Gaussian in the residuals between data and waveform (Finn 1992; Cutler & Flanagan 1994; Abac et al. 2026a). We use the standard formulation for this likelihood and the associated frequency-domain noise-weighted inner product (e.g., Veitch et al. 2015b; Thrane & Talbot 2019). Given the likelihood and a prior distribution $\pi(\boldsymbol{\theta})$, we compute the posterior distribution from Bayes' theorem

$$p(\boldsymbol{\theta}|d) = \frac{p(d|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{p(d)}\,, \tag{9}$$

where $p(d)$ is the normalizing evidence:

$$p(d) = \int p(d|\boldsymbol{\theta})\pi(\boldsymbol{\theta})\,d\boldsymbol{\theta}\,. \tag{10}$$

The evidence is not required in most approaches to PE, but can be used for model comparison.

In our baseline analyses the joint likelihood also includes additional parameters to account for the calibration state of the detector and the uncertainty in this calibration (Section 5.4). The full likelihood is evaluated coherently across the detector network by multiplying the individual interferometer likelihoods, under the assumption that the noise is independent in each.

5.2. *Preliminary PE*

After initial identification of a candidate by a search algorithm (Section 3), we carry out preliminary PE in order to understand the properties of the source and guide our analysis settings. We choose our initial settings and waveform model using the output of the search pipelines, e.g., the point estimates of the source masses. As needed, we carry out further analyses with modified priors or configuration settings, in some cases utilizing multiple waveform models to understand systematic modeling errors. If there are data-quality issues, this preliminary estimation is iterated alongside glitch-mitigation studies, as described in Section 4.

Preliminary PE also guides our choice of waveform models for inference (Section 2).

5.3. *Likelihood Calculation: Elements of the Inner Product*

The integration of the inner product that defines the likelihood is carried out in the frequency domain (Abac et al. 2026a; Veitch et al. 2015b; Thrane & Talbot 2019). Following our standard conventions (Abac et al. 2026a), given a discrete raw time-series data $d(t)$ with a duration T and sampling frequency $f_{\rm s}$, we first apply a FFT and divide by the sampling frequency to produce the discrete frequency series $\tilde{d}$. Similarly, time-domain waveforms are transformed into frequency series, while frequency-domain models can be evaluated directly for use in the inner product.

The amount of data analyzed for each candidate depends on the estimated mass of the system as determined by preliminary PE analyses, with lower mass signals in the sensitive band for longer periods of time. The segment of data used in PE is selected with the estimated time of coalescence set 2 s before the end of the segment, with the total duration set to powers of 2 ranging from 4 s up to 128 s such that the evolution of the signal starting from $f_{\rm low}$ is included. This choice sets the discrete frequency spacing $\Delta f = 1/T$ in the inner product. Before transforming time-domain data to the frequency domain to evaluate the likelihood, we apply a symmetric Tukey window function (Harris 1978) with a roll-on time that balances how much data in the segment is affected by the windowing and the impact of spectral leakage of the noise, especially at the lowest frequencies analyzed where the noise is steeply varying.

We compute the required integral from a lower frequency $f_{\rm low}$ to $f_{\rm high}$ and discard the data outside this range. Generally, $f_{\rm low}$ is set to 20 Hz, below which the detector noise steeply increases. The upper frequency $f_{\rm high}$ is chosen to be the Nyquist frequency, half the discrete sampling rate $f_{\rm s}$ of the data, with further modifications to account for power lost due to applying a low-pass Butterworth filter when downsampling the data to the desired sampling rate (Veitch et al. 2015a; Romero-Shaw et al. 2020b). The result is $f_{\rm high} = \alpha^{\text{roll-off}} f_{\rm s}/2$, with the particular choice of $\alpha^{\text{roll-off}} = 0.875$ selected to limit power loss to 1% (Abbott et al. 2023a). The data are downsampled to limit the computational cost of the likelihood evaluation, and $f_{\rm s}$ is chosen for each candidate to ensure that selected higher multipolar modes are resolved for candidates which coalesce in the sensitive frequency band of the detectors, or else set to a high enough value where noise dominates the signal power for low-mass candidates that coalesce out of the sensitive band. This upper limit is generally $f_{\rm s} = 4096$ Hz or $f_{\rm s} = 8192$ Hz.

A key ingredient in the likelihood evaluation is an estimate of the noise PSD for the analyzed data.

In order to mitigate the effects of the nonstationary nature of the instrumental noise, for PE results presented in GWTC-5.0, we use an on-source estimate for the PSD generated with the BAYESWAVE algorithm (Cornish & Littenberg 2015; Littenberg & Cornish 2015; Cornish et al. 2021; Gupta & Cornish 2024).

This approach uses Bayesian inference to construct a posterior distribution over PSD realizations, and we take the median PSD value at each frequency (Chatziioannou et al. 2019) for the point-estimate used in the inner product.

5.4. *Calibration Marginalization*

In order to account for the uncertain calibration of the strain data, we allow for independent frequency-dependent shifts of the amplitude and phase of the waveform in each interferometer (Farr et al. 2014). The frequency-dependent shifts are modeled using additional free parameters in our waveform model, which we call calibration parameters. We marginalize over the calibration parameters when reporting PE results (Abbott et al. 2016d).

The strain data are produced in near-real time using an initial calibration model, which may be subsequently updated in order to recalibrate the data (Abac et al. 2026c). For the LIGO detectors the methods used to calibrate the strain data provide a measured distribution of frequency-dependent correction factors $\tilde{\eta}_R(f)$ (Cahillane et al. 2017; Sun et al. 2020, 2021; Dartez et al. 2025). These factors correct the calibrated strain data d in each detector to the strain that would be measured with perfect calibration $d_\star$,

$$\tilde{d}_\star = \tilde{\eta}_R \tilde{d}\,. \tag{11}$$

For the Virgo detector the methods used for calibration give an estimate of the inverse of the correction factor $1/\tilde{\eta}_R$ (which has been unity through O3) and the corresponding frequency-dependent uncertainties (Accadia et al. 2014; Acernese et al. 2018, 2022).

The likelihood used in PE relies on knowledge of the noise properties of the imperfectly calibrated data d through the estimated PSD. The likelihood $p(d|\boldsymbol{\theta})$ is evaluated in the frequency domain using $\tilde{d}$,

$$\tilde{d} = \frac{1}{\tilde{\eta}_R}[\tilde{h}(\boldsymbol{\theta}) + \tilde{n}_\star] = \frac{\tilde{h}(\boldsymbol{\theta})}{\tilde{\eta}_R} + \tilde{n}\,, \tag{12}$$

where $\tilde{n}_\star$ is the true noise in the detector, assumed to be Gaussian, and $\tilde{n} = \tilde{n}_\star/\tilde{\eta}_R$ is the noise in the calibrated strain, also assumed to be Gaussian. The PSD describes the properties of $\tilde{n}$, so the correction factors $\tilde{\eta}_R$ must be applied to correct the GW waveform $\tilde{h}$ to account for imperfect calibration. The corrections $1/\tilde{\eta}_R$ are approximated by small frequency-dependent amplitude and phase corrections, which are modeled using splines (Farr et al. 2014). The values of these splines at fixed frequency nodes are the calibration parameters, and for PE we use Gaussian priors on these additional parameters. For LIGO data we use the median and standard deviation of the measured distribution of $\tilde{\eta}_R$ (Cahillane et al. 2017; Sun et al. 2020, 2021; Dartez et al. 2025) to set the mean and standard deviation of the priors over the calibration parameters describing $1/\tilde{\eta}_R$. For Virgo data, the priors on the calibration parameters through O3 are zero-mean Gaussians with standard deviations corresponding to the directly measured uncertainties in $1/\tilde{\eta}_R$. For O4, frequency-dependent error is also provided, even if it was mainly negligible since the error was corrected in low-latency to provide an (almost) unbiased strain.

Up to O3, the data from LHO and LLO were recalibrated prior to final analysis in order to approximately remove the

systematic miscalibration arising from the initial model. For PE of these candidates, we used final calibrated data as described in Table 4 of Abac et al. (2026c). Because of this recalibration, the errors and uncertainty on the calibration parameters are small, such that the priors for the calibration parameters are nearly centered at zero (corresponding to no calibration correction). However, in O4 the data in LHO, LLO and Virgo are no longer recalibrated in this way by default, except for an approximately two-month period in O4b when recalibration was needed (Abac et al. 2026c). Nevertheless, recalibration during PE can be effectively carried out using calibration priors centered away from zero. For this reason the priors associated with the new GW candidates presented in GWTC-5.0 may lie more than a standard deviation away from zero for GWTC-5.0 at various frequency points.

5.5. *Priors*

To estimate the posterior distribution, Equation (9), we must select an appropriate prior distribution over the binary parameters $\boldsymbol{\theta}$. Since our inferences are made over a catalog of CBC systems with a wide range of properties and using a variety of waveform models, our prior ranges are selected to be appropriate for each candidate. The priors chosen are agnostic and wide enough to cover the region of parameter space where the posteriors have support, while ensuring a reasonable amount of analysis time and accounting for the parameter ranges over which our waveform models are calibrated. Prior ranges (e.g., on the component masses) are selected during preliminary PE analysis and adjusted afterward as needed, for example to ensure that an arbitrary prior boundary does not affect the posterior.

For all candidates we use priors that are uniform over the (redshifted) detector-frame component masses $(1+z)m_i$, with boundaries in detector-frame total mass, detector-frame chirp mass, and mass ratio appropriate for each candidate and applied model (see Section 2). The priors are uniform on the spin magnitudes and isotropic in spin orientations. We use isotropic priors on the binary orientation, and priors that are uniform in comoving volume and comoving time so that the priors on the sky location are isotropic. For this we carry out sampling in a reference cosmology (Abac et al. 2026a) and reweight our final samples to priors appropriate for the cosmological model of Ade et al. (2016), which is a flat-ΛCDM model with $H_0 = 67.9\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ and $\Omega_\mathrm{m} = 0.3065$, following our approach in past versions of the GWTC.

5.6. *Sampling from the Posterior*

Due to the high dimensionality of the source parameters $\boldsymbol{\theta}$, brute-force evaluation of the posteriors is intractable. Instead we represent the posteriors for each candidate with discrete samples $\boldsymbol{\theta}_i$ which represent fair draws from the posterior distribution. The challenge then reduces to the problem of stochastically sampling from the posterior distribution $p(\boldsymbol{\theta}|d)$.

A number of algorithms exist to tackle this sampling challenge. These include nested sampling (Skilling 2006), a Monte Carlo technique to estimate the evidence and which produces samples as a byproduct, and Markov chain Monte Carlo (MCMC; Hastings 1970), which samples through random walks and forms the core of many sampling algorithms. The evidence can be computed with MCMC methods by e.g., thermodynamic integration (Goggans & Chi 2004; Littenberg & Cornish 2009) via parallel tempering (Earl & Deem 2005; van der Sluys et al. 2008; Veitch et al. 2015b). For many analyses, we use the DYNESTY (Speagle 2020) nested-sampling algorithm, and accessed through the BILBY package (Ashton et al. 2019; Romero-Shaw et al. 2020b) which includes a custom stepping method optimized for use on parallelized computing resources, specifically large-core-count CPUs. While for older results parameter inference was carried out using LALINFERENCE (Veitch et al. 2015b), which provides an implementation of MCMC and nested sampling, as well as tailored methods for proposing new sample points.

We also make use of additional, highly parallelized sampling techniques to tackle computationally expensive analyses. These include PARALLELBILBY (Smith et al. 2020) and the RIFT algorithm (Pankow et al. 2015; Lange et al. 2017; Wysocki et al. 2019). PARALLELBILBY is optimized to use distributed computing to perform nested sampling. RIFT is a highly parallel, iterative sampling algorithm that uses adaptive, grid-based explorations of the likelihood while marginalizing over the extrinsic parameters, followed by a final stage of Monte Carlo sampling of the intrinsic and extrinsic parameters.

For the new candidates presented in GWTC-5.0, we employ nested sampling in BILBY as well as the RIFT algorithm to perform sampling. Additionally, BILBYMCMC (Ashton & Talbot 2021), which is a parallel-tempered ensemble Metropolis–Hastings Monte Carlo sampling algorithm, and a new machine-learning based pipeline, DINGO (Green et al. 2020; Dax et al. 2021, 2023) trained using SEOBNRV5PHM, were used for cross-validation. DINGO employs neural posterior estimation to train a normalizing flow-based neural density estimator that directly approximates the posterior distribution. Once fully trained using simulated data, the model enables amortised inference, reducing parameter estimation times by several orders of magnitude relative to conventional stochastic samplers. A subsequent importance sampling step is used to correct for residual discrepancies between the learned posterior and the true likelihood.

To improve sampling performance, certain parameters in the waveform model can be marginalized out of the posterior during sampling. Depending on the waveform model employed and the GW candidate, any of the coalescence phase (Veitch et al. 2015b; Farr 2014), coalescence time (Farr 2014; Romero-Shaw et al. 2020b), or luminosity distance (Singer & Price 2016) have been marginalized over, as discussed in Section 5.9. This can be done analytically or numerically depending on the parameter and the details of the waveform model. When marginalization is carried out, the full posterior distributions are reconstructed in postprocessing (Thrane & Talbot 2019).

The method of marginalizing over calibration uncertainties varies depending on the sampling method used. For inferences carried out using RIFT, the marginalization over the calibration model parameters is carried out in post-processing, by reweighting the likelihood of samples produced first without utilizing the spline calibration model (Payne et al. 2020; Abbott et al. 2023a). Meanwhile, for PE using BILBY, PARALLELBILBY, and LALINFERENCE, the parameters of the calibration model are inferred alongside the intrinsic and extrinsic binary parameters during sampling.

A large number of PE analyses and steps within each analysis is required to produce the inferences included in GWTC-5.0. These analyses were managed using the ASIMOV software library (Williams et al. 2023). Postprocessing of the samples is carried out with the PESUMMARY software library (Hoy & Raymond 2021). This includes the computation of derived parameters such as the remnant properties following BBH merger and the evolution of the spin components to large binary separations as described in Section 2.4. These are implemented with routines in LALSUITE (LIGO–Virgo–KAGRA Collaboration 2018; Wette 2020) and used by PESUMMARY. For reweighting of samples, e.g., to our preferred cosmological model or to incorporate calibration marginalization with RIFT, either routines in PESUMMARY or simple custom routines (as in the case for new candidates in GWTC-5.0) are used and provided with our data release.

5.7. *Posterior Samples*

Following sampling, our measurements of the binary parameters $\boldsymbol{\theta}$ for each candidate are represented by a discrete set of parameter values $\boldsymbol{\theta}_i$ sampled from the posterior distribution. These samples can be used to compute expectation values over quantities of interest, e.g., through Monte Carlo integration. Marginalization is achieved by only considering the dependence of the samples on the quantities that are not marginalized out. Each sample includes the fifteen intrinsic and extrinsic parameters required to describe the GW strain from quasi-circular BBHs as well as the parameters required to describe the inferred calibration state. In addition to these a number of quantities can be derived from the intrinsic parameters of the binary using the methods described in Section 2.4. These include the final mass $M_{\rm f}$ and final dimensionless spin $\chi_{\rm f}$ of the remnant compact object following coalescence, the peak luminosity $\ell_{\rm peak}$ and total energy radiated $E_{\rm rad}$, as well as the final kick velocity of the remnant $v_{\rm k}$ when the relevant waveform model supports such estimates.

Some intrinsic properties of the binaries vary over the course of the binary coalescence, i.e., the orientations of the dimensionless spins $\boldsymbol{\chi}_1$ and $\boldsymbol{\chi}_2$ in precessing systems. Our initial inferences report these quantities at a reference frequency $f_{\rm ref}$, taken to be 20 Hz, corresponding to a variable reference time before coalescence. The spin tilt angles with respect to the orbital angular momentum, θ_1 and θ_2, in particular carry valuable information about the formation mechanism of the binary. These angles approach well-defined limits at infinite binary separation (Gerosa et al. 2015), where they can be used in population inferences (e.g., Mould & Gerosa 2022), and so we report the spin tilts and derived quantities such as $\chi_{\rm eff}$ and $\chi_{\rm p}$ in this limit. The derived spins are evolved from their reference values (Johnson-McDaniel et al. 2022b) using PN expressions for the precession-averaged dynamics (Gerosa et al. 2015; Chatziioannou et al. 2017) to formally infinite separations (cf. Gerosa et al. 2023). Wherever possible, the more accurate hybrid evolution using a combination of orbit-averaged precession followed by precession-average evolution is employed, but for a small number candidates only the computationally faster precession-averaging approach is used. We report these evolved spin quantities in our final samples for BBH candidates, and for NSBH candidates when analyzing them with waveform models that neglect the imprint of matter onto the GW signal.

In reporting and plotting the inferred parameters for each candidate, we marginalize out all but one or two parameters at a time, and report the median value and credible intervals (CIs) of the marginalized posteriors for those parameters. Unless otherwise noted, we report the median and 90% CI of each binary parameter, marginalizing over the others. Generally we use symmetric CIs for single parameters, so that 5% of the posterior density lies below the lower bound of the 90% CI and 5% of the density lies above the upper bound. In some instances where the posteriors have support near a physical prior boundary, the symmetric CI gives the appearance of excluding the boundary value even if there is high probability there. In such cases we may report the 90% highest posterior density (HPD), the smallest interval containing 90% of the posterior density. When plotting marginalized densities we make use of one- and two-dimensional KDE to produce continuous densities from samples. Our two-dimensional credible regions are constructed using the HPD method.

Our inferences on the location of GW candidates are available in two formats. The samples themselves represent our full posterior over the source parameters, including the sky location and distance of the detected GW sources. In addition, we provide three-dimensional localizations in the same format as the localizations included in our public GW alerts, by applying KDE to the samples. These localizations are created using LIGO.SKYMAP (Singer & Price 2016; Singer et al. 2016; Singer et al. 2016), which includes the BAYESTAR package used to localize GW candidates from modeled online and offline searches.

5.8. *Approximate Spatial Localization*

We do not carry out full PE for all candidates which meet the criteria for inclusion in GWTC-5.0 (Section 3.6). Our criteria for full PE analysis have evolved with GWTC versions (Section 5.9) due to the increased number of detection candidates and resource limitations. However, we provide approximate spatial localization for all candidates in GWTC-5.0 to enable multimessenger analyses of large samples of weak GW candidates using the same methods as for our public GW alerts. Such methods are computationally much cheaper than

full PE, and already integrated into the candidate management system. In particular, for candidates produced by modeled CBC searches, the approximate localization is carried out with BAYESTAR (Singer & Price 2016; Singer et al. 2016; Singer et al. 2016). For candidates produced by cWB-BBH, the localization method is directly part of cWB and described in Section 3.1.

5.9. *Analysis Settings and Details*

Here we describe the particular analysis settings which vary across the inferences carried out for each candidate in GWTC-5.0.

5.9.1. *Candidates Found in O4*

For the new candidates from O4a and O4b presented in GWTC-5.0, we perform full Bayesian PE on a high-purity subset of the candidates (Section 3.6), namely those with FARs $< 1\,\mathrm{yr}^{-1}$ and $p_{\mathrm{astro}} > 0.5$ and not determined to be of instrumental origin (Section 4).

For this subset, the analysis settings depend on the nature of the binary as inferred using preliminary PE and confirmed in our final analysis. As described in Abac et al. (2026b) these candidates include a large number of BBHs but no significant NSBH or BNS candidates.

Based on preliminary PE and any input from data validation (cf. Section 4), we categorize the signals and determine an appropriate prior, waveform model, and data segment to analyze (in all cases, the data product is based on the `GDS-CALIB_STRAIN_CLEAN_AR` channel, as described in Abac et al. (2026c)). For nearly all cases, the durations are selected so that the $\ell = 3$, $|m| = 3$ higher harmonic is fully captured within the frequency range integrated over in our likelihood. In higher-mass systems $\ell = 4$ modes may make a non-negligible contribution in our sensitive frequency band (e.g., Mills & Fairhurst 2021), and in such cases we set the starting frequency so that the $\ell = 4$, $|m| = 4$ mode begins in the range of integration. Specifically, we do so for events whose source properties have support for redshifted mass $(1+z)M > 200M_{\odot}$. In each case it is possible that at the starting frequency of the waveform even higher harmonics will then begin in band, which may cause aliasing in the case of time-domain waveform models. Due to the small amplitude of such higher harmonics relative to the lower-frequency multipole moments, the effect is expected to be negligible. For all events, we adopt the additional prior on the binary parameters during sampling such that the $\ell = 3$ mode is resolved in band, given the chosen sampling rate. This only modifies the prior for a small number of candidates, and we have checked that in these cases, the prior makes no discernible difference in the final posterior samples.

For all candidates the Tukey window applied before transforming time-domain data to the frequency domain has a roll-on of 1 s. This is longer than the 0.4 s roll-on used in prior PE analyses. The longer window is chosen to reduce spectral leakage, because for the first time the instrumental noise near f_{low} is sufficiently low that the small amount of leakage from even lower frequencies into frequency range integrated over in our likelihood made noticeable impact on our analyses.

The analysis of the initial PE results to determine our final settings is carried out using the PECONFIGURATOR software package. The resulting information is also used to determine the appropriate waveform models to use to capture systematic uncertainties across waveform models.

For BBH candidates, we use both the phenomenological models IMRPHENOMXPHM_SPINTAYLOR (García-Quirós et al. 2020; Pratten et al. 2020a; Colleoni et al. 2025a) and IMRPHENOMXPNR (Hamilton et al. 2026) and the effective one body model SEOBNRV5PHM (Khalil et al. 2023; Pompili et al. 2023; Ramos-Buades et al. 2023; Estellés et al. 2026) for all candidates. In addition, we use NRSUR7DQ4 (Varma et al. 2019) for those candidates whose parameters lie within the range of total mass and mass ratio values supported by the model, as determined by preliminary PE. In the case of NRSUR7DQ4 we use a fixed duration of $10000(G/c^3)M(1+z)$ for the waveform model when evaluating the likelihood, with $M(1+z)$ the total detector-frame mass of the source.

For all new candidates presented in GWTC-5.0, the luminosity distance is marginalized over during sampling. For sampling we use a distance prior uniform in comoving volume and comoving time with the default cosmological model of the ASTROPY software package, a flat-Λ cold dark matter (CDM) cosmology with $H_0 = 67.66\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ and $\Omega_{\mathrm{m}} = 0.30966$. The cosmology assumed during sampling is not the default cosmology we present our results in. Instead, the final samples are reweighted to our preferred cosmology (Ade et al. 2016) as described in Section 5.5.

The analysis of the candidates from O4a first presented in GWTC-4.0 used the same methods and settings as described above, except for the choice of waveform models as per Table 1, and the use of the DYNESTY sampling algorithm through the BILBY package for all PE results.

5.9.2. *Candidates found in O1, O2, O3*

As a cumulative catalog, GWTC-5.0 includes candidates detected during the first three observing runs and the first part of the fourth observing run of the advanced detector network, as discussed in Section 3.6. For GW candidates identified in O1, O2, O3 and O4a we have previously performed full PE for the subset of these candidates identified with $p_{\mathrm{astro}} > 0.5$ plus the NSBH candidate GW200105_162426 (Abbott et al. 2021b), and our PE results for these candidates remain the same in GWTC-5.0. These inferences were first presented in GWTC-2.1 (Abbott et al. 2024, for candidates found in O1, O2, and O3a) and GWTC-3.0 (Abbott et al. 2023a, for candidates found in O3b). The exception is the BNS candidate GW170817, for which the PE results were not updated in GWTC-2.1, and so remain the same as those presented in GWTC-1.0 (Abbott et al. 2019a). Both sets of analyses from GWTC-2.1 and GWTC-3.0 used the same methods and settings which we now summarize.

For each candidate, PSDs were generated using BAYESWAVE (Cornish & Littenberg 2015; Littenberg &

Cornish 2015; Cornish et al. 2021; Gupta & Cornish 2024) as described in Section 5.3.

For these candidates the roll-on of Tukey window applied to the time domain-data before transforming to the frequency domain was $0.4\,\mathrm{s}$. Multiple waveform models were used in each case in order to understand systematic uncertainties in the inferences.

For BBH candidates, PE was carried out using two waveform models which include higher harmonics and the effects of orbital precession. The phenomenological model IMRPhenomXPHM (García-Quirós et al. 2020; Pratten et al. 2020a) was used with a multiscale prescription for the precession dynamics (Chatziioannou et al. 2017) and sampling was carried out with Bilby. The EOB model SEOBNRv4PHM (Ossokine et al. 2020) was also used for each BBH candidate, using RIFT to carry out PE.

For the NSBH candidates, multiple waveform models were also used for PE. Our baseline results are drawn from the same models as for BBH candidates, namely IMRPhenomXPHM and SEOBNRv4PHM. These models include higher harmonics and precession but neglect matter effects such as the tidal deformation or disruption of the NS on the waveform. This is because such effects are negligible for the NSBH candidates identified in O3b (Abbott et al. 2021b). In order to assess the importance of matter effects on the GW signal, models which include the effects to tidal deformation and models which are specifically tailored for NSBH systems were also used for inference. These were IMRPhenomNSBH and SEOBNRv4_ROM_NRTidalv2_NSBH. In addition, IMRPhenomPv2_NRTidalv2 is employed to approximately incorporate both matter and precession effects. There were no new NSBH candidates detected in O4b.

Up to GWTC-4.0, two BNS candidates have been identified satisfying the PE threshold, GW170817 and GW190425. Multiple models are used for PE of GW170817 (Abbott et al. 2019a). GW170817 was analyzed with three frequency-domain models using LALInference: TaylorF2 including tidal effects (Sturani 2015; Isoyama et al. 2020; Flanagan & Hinderer 2008; Vines et al. 2011), IMRPhenomPv2_NRTidal, and SEOBNRv4_ROM_NRTidal, and with two time-domain models using RIFT: SEOBNRv4T (Hinderer et al. 2016; Steinhoff et al. 2016), and TEOBResumS (Nagar et al. 2018). For this candidate we marginalize over the phase analytically for models that assume spins aligned with the orbital plane. As with our other analyses on candidates including NSs, multiple analyses are carried out allowing only for relatively small spin magnitudes $\chi_i \leq 0.05$ and allowing for spin magnitudes up to the maximum allowed for a given model, as large as $\chi_i \leq 0.99$, which allows us to infer an upper bound on the component spin magnitudes from the data without prior constraints (Abbott et al. 2019c). The second BNS candidate, GW190425 (Abbott et al. 2020b), was detected during O3a and its PE results updated in GWTC-2.1 (Abbott et al. 2024). For this candidate we present results using the precessing, tidal approximant IMRPhenomPv2_NRTidal using Dynesty through Bilby, and both relatively low- and high-spin prior limits on the component spins. To accelerate inference we employ reduced-order quadrature.

When using Bilby or ParallelBilby to analyze candidates from the first three observing runs, the posteriors are marginalized over luminosity distance and geocenter time, with the exception of the Bilby analysis of GW190425 which only used distance marginalization. For these candidates, initial sampling was carried out with a distance prior uniform in Euclidean volume. During postprocessing, the posterior samples were then reweighted to the cosmological model described in Section 5.5 (Ade et al. 2016).

5.9.3. *Calibration Prior Settings for Candidates from O1, O2 and O3*

We marginalized over the calibration uncertainties when producing the PE results for O1, O2, and O3, as discussed in Section 5.4. Due to an error in implementation, an incorrect prior on the calibration parameters for LHO and LLO was used for these results. The priors were set using the median and CIs of $\tilde{\eta}_R$ from Equation (11) for data from the LIGO detectors, rather than those of $1/\tilde{\eta}_R$ as required by the method. The inconsistency was discovered in the analysis of GW240925_005809, where the LHO calibration was inferred directly from the signal (Abac et al. 2026f). In the limit of small calibration uncertainties, this amounts to a sign error in the means of the Gaussian priors for the calibration parameters. In the case that the means of the priors are zero the error has no effect. However, the calibration uncertainties for LHO and LLO have nonzero means. Meanwhile, the priors on the calibration parameters were set correctly for Virgo data.

For candidates detected in O1, O2, and O3 the means are generally small relative to the standard deviation of the priors on the calibration parameters. Further the absolute sizes of the standard deviations of these parameters are small in the sensitive band of the detectors, of the order of a few percent in amplitude and a few degrees in phase for the two LIGO interferometers (Abbott et al. 2019a, 2021a, 2024, 2023a), and so *a priori* the impact of this error on our inferences is expected to be small. We have verified this expectation through preliminary re-analysis of the potentially impacted candidates, carried out by repeating PE with corrected priors and by reweighting the likelihood values of existing PE samples in order to correct for the erroneous calibration priors (Baka et al. 2025). None of the scientific conclusions reported in previous studies are affected by the error. In addition, the error does not impact the significance of any of our candidates, since the GW searches described in Section 3 do not incorporate the effect of uncertain strain calibration.

The typical change in the posteriors following preliminary re-analysis is within statistical sampling error, as quantified by the Jensen–Shannon divergence (Lin 1991) between one-dimensional marginal distributions before and after correction (cf. Romero-Shaw et al. 2020b; Abbott et al. 2021a, Appendix A). As a particular case, the localization of the source of GW170817 receives only a small correction when re-

analyzed, and its association with the electromagnetic counterpart emission from AT 2017gfo (Abbott et al. 2017b) is unaffected by the error.

Another case is GW150914, which displays visible differences in the sky location posteriors when re-analyzed (although the bounds of the 90% CI of the right ascension and declination remain nearly unchanged). Meanwhile, our inferences of the intrinsic parameters of GW150914 remain unchanged to within sampling errors.

While the impact on individual candidates is small, a possible concern is that this error can bias analyses that aggregate data from multiple candidates, such as population studies and cosmological inferences. We are investigating the impact of the calibration marginalization error on these analyses, but the effects are expected to be negligible compared to other sources of systematic error. This error does not impact the PE of new candidates observed in O4b (Abac et al. 2026b).

5.10. *On the Likelihood used for Inference in O1, O2 and O3*

During the production stage for O4a, we discovered a normalization error in the likelihood used for the inference codes BAYESWAVE, BILBY, LALINFERENCE, PARALLELBILBY, and RIFT. The error arises due to the incorrect application of a window factor that was intended to account for the power lost in the noise residuals due to the Tukey window applied to the data before transforming them to the frequency domain. This error causes the likelihood to be overly constrained by a factor depending on the window function used to mitigate spectral leakage. The incorrect likelihood $\hat{p}(d|\boldsymbol{\theta})$ is related to the correct likelihood $p(d|\boldsymbol{\theta})$ via the average power in the Tukey window $\hat{p}(d|\boldsymbol{\theta}) = p(d|\boldsymbol{\theta})^{\beta}$ with

$$\beta = \left(1 - \frac{5T_{\rm w}}{4T}\right)^{-1}, \tag{13}$$

where $T_{\rm w}$ is the roll-off time of the Tukey window to one side, and T is the segment duration. The window factor is applied when computing the PSD via standard methods, but should not be applied to the data when the signal has support only where the window function is unity, as is the case in our analyses. Although accounting for the windowing using a single window factor in PSD estimation is still an approximation when computing the likelihood, multiple investigations and the use of probability–probability tests (Veitch et al. 2015b; Romero-Shaw et al. 2020b) have confirmed its accuracy. Further, these tests have confirmed the normalization error in our inference codes and the correctness of our updated likelihood. More details on the error and validation of the updated likelihood can be found in Talbot et al. (2025).

For the new candidates from O4, we used the correct likelihood $p(d|\boldsymbol{\theta})$. Comparisons with preliminary results for O4a generated with the incorrect likelihood $\hat{p}(d|\boldsymbol{\theta})$ show that the differences in the posteriors between the two methods are small, but correcting the error systematically widens the posterior distributions. This is the expected behavior, since the effect of the error corresponds to systematically overestimating the SNR inferred from PE.

While the error is not present in the source properties reported in GWTC-4.0 (Abac et al. 2025a), due to resource constraints, posteriors from earlier analyses that included this error were used in the waveform consistency tests (Abac et al. 2025b) for the new candidates from O4a. The error in the likelihood remains present in the BAYESWAVE analyses used in those tests. For the worst cases, $T_{\rm w} = 1\,{\rm s}$ and $T = 4\,{\rm s}$, yielding $\beta = 1.45$. This corresponds to overestimated SNRs from PE by a factor 1.21. The impact is less for lower-mass candidates whose durations are larger, so that for BBH candidates with $T_{\rm w} = 1\,{\rm s}$ and $T = 8\,{\rm s}$, the factor is $\beta = 1.19$, and the SNRs are overestimated by a factor of 1.09. To check the impact of this error on the waveform consistency tests, we reran the consistency tests for a subset of events using the correct posterior distributions, and find no significant changes in the results reported in Abac et al. (2025a).

For O4a candidates, we released our final posteriors generated using the correct likelihood $p(d|\boldsymbol{\theta})$; posteriors generated using the incorrect likelihood $\hat{p}(d|\boldsymbol{\theta})$; and a set of reweighted posteriors that correct the error by applying rejection sampling with replacement to posteriors generated with $\hat{p}(d|\boldsymbol{\theta})$, with an acceptance ratio given by the ratio of the correct to incorrect likelihood.

The error does affect previous analyses from O1 through O3, but the impact is smaller than for O4 because of the smaller Tukey window applied when performing PE for candidates from those observing runs, and we have not corrected these past inferences. For results from these previous runs, the worst cases have $T_{\rm w} = 0.4\,{\rm s}$ and $T = 4\,{\rm s}$, yielding $\beta = 1.14$. In this case, the SNRs are overestimated by a factor of 1.07.

This normalization error in the likelihood does not impact the PE results for the new candidates presented in GWTC-5.0, since the correct likelihood $p(d|\boldsymbol{\theta})$ was used for those analyses.

6. WAVEFORM CONSISTENCY TESTS

As we have seen in the earlier sections, for our modeled analyses we use several waveforms based on different approximations. Events like GW190521 (Abbott et al. 2020d), GW200105_042309 (Abbott et al. 2021b), and GW231123_135430 (Abac et al. 2025d) highlight the importance of continuously checking the validity of our assumptions, in order to ensure the reliability of astrophysical interpretations. To this end, we perform *waveform consistency tests* that compare different waveform reconstruction techniques to assess their agreement and identify any unexpected mismatch between the reconstructed signals. Potential mismatches can arise from assumptions at the pipeline level such as the noise properties, the effects of the subtraction of glitches from the data stream, and inadequate waveform choice. Mismatches can also originate from physical or astrophysical properties, such as deviations from GR, the presence of eccentricity unaccounted for in waveforms, and gravitational lensing-induced deformations of the measured waveforms.

Waveform reconstruction techniques can be minimally modeled (as seen for cWB in Section 3.1) or template based (Sections 3 and 5). Minimally-modeled techniques use time–frequency representations, based on wavelets or other time–frequency techniques, to identify coherent features in the data of a network of multiple detectors. This generic approach enables the discovery of unexpected phenomena but does not provide a direct mapping between the reconstructed waveform and the source's physical properties, such as the masses and spins of a binary system and the distance to the source.

To evaluate the consistency between template-based and minimally-modeled reconstructions, we implement a systematic injection study (Abbott et al. 2019a, 2021a, 2023a; Salemi et al. 2019; Ghonge et al. 2020; Johnson-McDaniel et al. 2022a). This involves injecting CBC waveform samples from the posterior parameter distributions (Section 5) into detector data at times near but distinct from the candidate (*off-source* injections). These injections are then reconstructed using minimally-modeled methods. By comparing the reconstructed waveforms from off-source injections (w_i) with the reconstruction of the candidate ($\hat{w}$), we can assess how well the CBC PE posteriors align with the minimally-modeled reconstruction. We quantify the agreement between waveform reconstructions using the *overlap*, defined as

$$\mathcal{O}(h_1, h_2) = \frac{\langle h_1 | h_2 \rangle}{\sqrt{\langle h_1 | h_1 \rangle \langle h_2 | h_2 \rangle}}\,, \tag{14}$$

where h_1 and h_2 are the waveforms being compared and $\langle \cdot | \cdot \rangle$ denotes the noise-weighted inner product (Abac et al. 2026a, Appendix B). The overlap, $\mathcal{O}(h_1, h_2)$, is bounded between -1 and $+1$.

For each candidate, we compute the empirical distribution $\mathcal{O}$ of *off-source* overlaps $\mathcal{O}(w_i, h_i)$ between injected waveforms and their minimally-modeled reconstructions. Next, we use the *on-source* overlap $\mathcal{O}(\hat{w}, h_{\mathrm{maxL}})$ between the maximum-likelihood posterior sample and the actual minimally-modeled reconstruction to compute a p-value on the left tail of the overlap distribution (i.e., for low overlap values).

This comparison has an inherent asymmetry: for off-source cases, we calculate the overlap between known waveforms and their minimally-modeled reconstructions, while for on-source cases, we compare the maximum-likelihood template from PE with the minimally-modeled reconstruction. This asymmetry typically results in off-source overlaps being systematically lower than what would be expected from a true null distribution (Abbott et al. 2023a). Consequently, the derived p-values are conservative by construction. Despite this limitation, the on-source p-value remains a useful indicator for identifying unexpected signal features, though with reduced statistical power compared to an unbiased test.

To ensure reliable waveform consistency tests, we selected all O4b candidates that meet four key criteria: (i) data availability from both LIGO detectors, (ii) presence of PE results using the IMRPhenomXPHM_SpinTaylor waveform family, (iii) detector-frame chirp mass $(1+z)\mathcal{M} > 15 M_\odot$, and (iv) a network SNR > 10. These criteria reflect that minimally-modeled methods require multi-detector data, and that they are particularly effective for high-mass CBCs with high SNR signals.

The candidates satisfying these criteria are analyzed in a companion paper (Abac et al. 2026b) using three minimally-modeled waveform reconstruction methods: BayesWave (Cornish & Littenberg 2015; Cornish et al. 2021; Ghonge et al. 2020; Gupta & Cornish 2024), and two configurations of the cWB pipeline. The two cWB configurations are cWB-2G (Klimenko et al. 2016; Drago et al. 2021) which employs the WDM wavelet transform and an excess power statistic to identify coherent features, and cWB-BBH (Klimenko 2022), also described in Section 3.1, which uses the WaveScan transform and a cross-power statistic. Both BayesWave and cWB-2G are designed for generic GW transients, whereas cWB-BBH is optimized for CBC signals, also adopting specialized frequency bands and time–frequency resolutions.

For each selected candidate, our injection campaign used for

- BayesWave approximately 200 random posterior samples injected within ± 8192 s of the candidate time,
- cWB-2G and cWB-BBH approximately 2000 random posterior samples injected across a 2–3 week period surrounding the candidate time.

This analysis provides a robust statistical baseline to quantify the degree of agreement between template-based and minimally-modeled reconstructions for a relatively large subset of the candidates from GWTC-5.0. By comparing the on-source overlap to the distribution from off-source injections, we can verify their consistency with the corresponding CBC reconstructions, and possibly identify those that may exhibit unexpected features.

7. DATA MANAGEMENT

The workflow of data analyses described in this paper outlines a complex chain of disparate analyses required to find and characterize GW transients (see Figure 1). In addition to the internal complexity of each analysis, coordinating each stage of this analysis and effectively tracking that input and output data is a significant challenge. This challenge grows with the number of GW candidates observed, necessitating the development of tools to manage these tasks with little to no human intervention. For O4, this development included the augmentation of existing infrastructure and the development of new software packages including cbcflow and the catalog data-product pipeline.

The online and offline analysis results from the search pipelines described in Section 3 are stored in GraceDB (Moe et al. 2014). For offline results, GraceDB is also used to generate the final candidate list for the catalog. To facilitate the tracking of these offline results, GraceDB has been augmented to provide version-controlled snapshots of the state of the catalog during the progression of the offline analyses.

As for catalogs since GWTC-2.0, the PE analyses described in Section 5 are managed with the ASIMOV software package (Williams et al. 2023). ASIMOV ingests data-quality recommendations and preliminary PE analyses, and uses this to automatically determine appropriate configuration for PE analyses, as well as automating the production of PSD estimates. Once all desired PE is complete, ASIMOV packages results into a standard format including all inputs and configuration required for reproduction using PESUMMARY (Hoy & Raymond 2021) .

The CBCFLOW (Ashton et al. 2022) software package is used to manage the flow of data between the various analyses and to track metadata about each stage of the analysis. A monitor process fetches search metadata from GRACEDB and preliminary PE results from a shared directory on the Caltech computing cluster. Data-quality information about candidates is updated following the studies described in Section 4. ASIMOV reads CBCFLOW for search, data quality, and preliminary-PE metadata, which is used to configure production analyses. Upon completion of those analyses, metadata about them are written back to CBCFLOW. All other downstream analyses such as searches for lensed pairs of candidates (Abac et al. 2026d) and tests of GR (Abac et al. 2026e) also utilize CBCFLOW to track their progress and results.

Finally, the release data product is assembled by a collection of scripts. These read CBCFLOW to identify the preferred search and PE results, as well as tracking their finalization status. These are then collated into the data product itself, and tables of summary information are generated to facilitate downstream use.

8. CONCLUSION

Leading on from the introduction presented in Abac et al. (2026a), this article describes the analysis methods used to transform the interferometric strain data from the LVK detectors into version 5.0 of the GWTC; the results of these analyses are presented in Abac et al. (2026b). We began in Section 2 with a description of the waveform models used to describe the GW signals from CBCs involving BHs and NSs. In Section 3, we described the search methods we use to filter the strain data to identify candidate transient GW signals (sensitive to both CBC sources as well as minimally-modeled bursts of GW radiation), and how these candidates are then ranked to identify the most significant ones. Likely GW candidates are then studied to understand the quality of data and identify any transient non-Gaussian noise that may bias later analyses (see Section 4). Section 5 described how a subset of likely GW candidates are then characterized using computational Bayesian inference methods to estimate the parameters of the GW source, such as the masses and spins of the compact objects involved. To check the assumptions made in the PE analysis, we described in Section 6 how waveform consistency tests compare different waveform reconstruction techniques to assess their agreement and identify any unexpected features in the reconstructed signals. Finally, in Section 7, we described the data management and workflow tools used to coordinate the various analyses and track the input and output data.

The methods described in this work are continually developed to improve the capabilities of the LVK network to detect and characterize GW transients. Specifically, as the detectors evolve towards higher sensitivity, and more detectors are added to the network, we continue to see a greater number of signals and more high-fidelity signals at larger SNR (Abbott et al. 2020). This necessitates efficiency improvements across all the methods to avoid computational bottlenecks and to ensure that the GW transient candidates can be processed in a timely manner. Moreover, the increased SNR requires refinements to the waveform models to reduce systematic biases as much as possible. The evolution of the detectors also necessitates continual improvements to the data analysis methods to ensure optimal data processing (e.g., as the noise floor of the detectors lowers, new glitch classes may become relevant, requiring development to searches and data-quality studies). Finally, as the time–volume explored increases, we hope to further expand the range of astrophysical sources that can be detected and characterized, which in turn requires the development of new models, search, and PE methods.

Data Availability: The data products generated by the methods described within this work are openly available in the GWTC-5.0 online catalog, which is hosted at https://gwosc.org/GWTC-5.0, documented further in Abac et al. (2026c) and the results described in Abac et al. (2026b).

ACKNOWLEDGEMENTS

This material is based upon work supported by NSF's LIGO Laboratory, which is a major facility fully funded by the National Science Foundation. The authors also gratefully acknowledge the support of the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO 600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS), and the Netherlands Organization for Scientific Research (NWO) for the construction and operation of the Virgo detector and the creation and support of the EGO consortium. The authors also gratefully acknowledge research support from these agencies, as well as by the Council of Scientific and Industrial Research of India; the Department of Science and Technology, India; the Science & Engineering Research Board (SERB), India; the Ministry of Human Resource Development, India; the Spanish Agencia Estatal de Investigación (AEI); the Spanish Ministerio de Ciencia, Innovación y Universidades; the European Union NextGenerationEU/PRTR (PRTR-C17.I1); the ICSC—CentroNazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by the European Union NextGenerationEU; the Comuni-

tat Autonòma de les Illes Balears through the Conselleria d'Educació i Universitats; the Conselleria d'Innovació, Universitats, Ciència i Societat Digital de la Generalitat Valenciana and the CERCA Programme Generalitat de Catalunya, Spain; the Polish National Agency for Academic Exchange; the National Science Centre of Poland and the European Union—European Regional Development Fund; the Foundation for Polish Science (FNP); the Polish Ministry of Science and Higher Education; the Swiss National Science Foundation (SNSF); the Russian Science Foundation; the European Commission; the European Social Funds (ESF); the European Regional Development Funds (ERDF); the Royal Society; the Scottish Funding Council; the Scottish Universities Physics Alliance; the Hungarian Scientific Research Fund (OTKA); the French Lyon Institute of Origins (LIO); the Belgian Fonds de la Recherche Scientifique (FRS-FNRS), Actions de Recherche Concertées (ARC) and Fonds Wetenschappelijk Onderzoek—Vlaanderen (FWO), Belgium; the Paris Île-de-France Region; the National Research, Development and Innovation Office of Hungary (NKFIH); the National Research Foundation of Korea; the Natural Sciences and Engineering Research Council of Canada (NSERC); the Canadian Foundation for Innovation (CFI); the Brazilian Ministry of Science, Technology, and Innovations; the International Center for Theoretical Physics South American Institute for Fundamental Research (ICTP-SAIFR); the Research Grants Council of Hong Kong; the National Natural Science Foundation of China (NSFC); the Israel Science Foundation (ISF); the US–Israel Binational Science Fund (BSF); the Leverhulme Trust; the Research Corporation; the National Science and Technology Council (NSTC), Taiwan; the United States Department of Energy; and the Kavli Foundation. The authors gratefully acknowledge the support of the NSF, STFC, INFN, and CNRS for provision of computational resources.

This work was supported by MEXT; the JSPS Leading-edge Research Infrastructure Program, JSPS Grant-in-Aid for Specially Promoted Research 26000005, JSPS Grant-in-Aid for Scientific Research on Innovative Areas 2402: 24103006, 24103005, and 2905: JP17H06358, JP17H06361, and JP17H06364, JSPS Core-to-Core Program A. Advanced Research Networks, JSPS Grants-in-Aid for Scientific Research (S) 17H06133 and 20H05639, JSPS Grant-in-Aid for Transformative Research Areas (A) 20A203: JP20H05854; the joint research program of the Institute for Cosmic Ray Research, University of Tokyo; the National Research Foundation (NRF); the Computing Infrastructure Project of the Global Science experimental Data hub Center (GSDC) at KISTI; the Korea Astronomy and Space Science Institute (KASI); the Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS), the AS Grid Center (ASGC), and the National Science and Technology Council (NSTC) in Taiwan under grants including the Science Vanguard Research Program; the Advanced Technology Center (ATC) of NAOJ; and the Mechanical Engineering Center of KEK.

Additional acknowledgments for support of individual authors may be found in the following document: https://dcc.ligo.org/LIGO-M2300033/public. For the purpose of open access, the authors have applied a Creative Commons Attribution (CC BY) license to any Author Accepted Manuscript version arising. We request that citations to this article use "A. G. Abac et al. (LIGO-Virgo-KAGRA Collaboration), ..." or similar phrasing, depending on journal convention.

Software: Plots were prepared with MATPLOTLIB (Hunter 2007), NUMPY (Harris et al. 2020), and TIKZ (Tantau 2023).

REFERENCES

Aasi, J., et al. 2013, Phys. Rev. D, 87, 022002, doi: 10.1103/PhysRevD.87.022002

Aasi, J., Abbott, B. P., Abbott, R., et al. 2015, Classical and Quantum Gravity, 32, 074001, doi: 10.1088/0264-9381/32/7/074001

Abac, A. G., Abouelfettouh, I., Acernese, F., et al. 2026a, To be published in this issue. https://dcc.ligo.org/LIGO-P2500701/public

—. 2026b, To be published in this issue. https://dcc.ligo.org/LIGO-P2600152/public

—. 2026c, To be published in this issue. https://dcc.ligo.org/LIGO-P2600085/public

—. 2026d, To be published in this issue

—. 2026e, To be published in this issue

Abac, A. G., Abouelfettouh, I., Acernese, F., et al. 2025a, arXiv e-prints, arXiv:2508.18082, doi: 10.48550/arXiv.2508.18082

—. 2025b, arXiv e-prints, arXiv:2508.18081, doi: 10.48550/arXiv.2508.18081

—. 2025c, ApJL, 995, L18, doi: 10.3847/2041-8213/ae0c06

—. 2025d, ApJL, 993, L25, doi: 10.3847/2041-8213/ae0c9c

Abac, A. G., Abouelfettouh, I., Acernese, F., et al. 2026f, Phys. Rev. Lett., doi: 10.1103/gzrj-mwv3

Abadie, J., Abbott, B. P., Abbott, R., et al. 2012, ApJ, 760, 12, doi: 10.1088/0004-637X/760/1/12

Abadie, J., et al. 2012, Phys. Rev. D, 85, 082002, doi: 10.1103/PhysRevD.85.082002

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, ApJL, 833, L1, doi: 10.3847/2041-8205/833/1/L1

—. 2016b, Physical Review X, 6, 041015, doi: 10.1103/PhysRevX.6.041015

—. 2016c, ApJS, 227, 14, doi: 10.3847/0067-0049/227/2/14

—. 2016d, PhRvL, 116, 241102, doi: 10.1103/PhysRevLett.116.241102
—. 2017a, PhRvL, 119, 161101, doi: 10.1103/PhysRevLett.119.161101
—. 2017b, ApJL, 848, L12, doi: 10.3847/2041-8213/aa91c9
—. 2019a, Physical Review X, 9, 031040, doi: 10.1103/PhysRevX.9.031040
—. 2019b, ApJ, 875, 161, doi: 10.3847/1538-4357/ab0e8f
—. 2019c, Physical Review X, 9, 011001, doi: 10.1103/PhysRevX.9.011001
Abbott, B. P., et al. 2020, Living Rev. Rel., 23, 3, doi: 10.1007/s41114-020-00026-9
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020a, Classical and Quantum Gravity, 37, 055002, doi: 10.1088/1361-6382/ab685e
—. 2020b, ApJL, 892, L3, doi: 10.3847/2041-8213/ab75f5
Abbott, R., Abbott, T. D., Abraham, S., et al. 2020c, PhRvD, 102, 043015, doi: 10.1103/PhysRevD.102.043015
—. 2020d, PhRvL, 125, 101102, doi: 10.1103/PhysRevLett.125.101102
—. 2021a, Physical Review X, 11, 021053, doi: 10.1103/PhysRevX.11.021053
—. 2021b, ApJL, 915, L5, doi: 10.3847/2041-8213/ac082e
Abbott, R., Abbott, T. D., Acernese, F., et al. 2023a, Physical Review X, 13, 041039, doi: 10.1103/PhysRevX.13.041039
—. 2023b, Physical Review X, 13, 011048, doi: 10.1103/PhysRevX.13.011048
—. 2024, PhRvD, 109, 022001, doi: 10.1103/PhysRevD.109.022001
Accadia, T., Acernese, F., Agathos, M., et al. 2014, Classical and Quantum Gravity, 31, 165013, doi: 10.1088/0264-9381/31/16/165013
Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, Classical and Quantum Gravity, 32, 024001, doi: 10.1088/0264-9381/32/2/024001
Acernese, F., Adams, T., Agatsuma, K., et al. 2018, Classical and Quantum Gravity, 35, 205004, doi: 10.1088/1361-6382/aadf1a
Acernese, F., Agathos, M., Ain, A., et al. 2022, Classical and Quantum Gravity, 39, 045006, doi: 10.1088/1361-6382/ac3c8e
—. 2023a, Classical and Quantum Gravity, 40, 185006, doi: 10.1088/1361-6382/acd92d
—. 2023b, Classical and Quantum Gravity, 40, 185005, doi: 10.1088/1361-6382/acdf36
Adams, T., Buskulic, D., Germain, V., et al. 2016, Classical and Quantum Gravity, 33, 175012, doi: 10.1088/0264-9381/33/17/175012
Ade, P. A. R., Aghanim, N., Arnaud, M., et al. 2016, A&A, 594, A13, doi: 10.1051/0004-6361/201525830
Agathos, M., Meidam, J., Del Pozzo, W., et al. 2015, Phys. Rev. D, 92, 023012, doi: 10.1103/PhysRevD.92.023012
Ajith, P., et al. 2007, Class. Quant. Grav., 24, S689, doi: 10.1088/0264-9381/24/19/S31
Ajith, P., Hannam, M., Husa, S., et al. 2011, PhRvL, 106, 241101, doi: 10.1103/PhysRevLett.106.241101
Akcay, S., Bernuzzi, S., Messina, F., et al. 2019, Phys. Rev. D, 99, 044051, doi: 10.1103/PhysRevD.99.044051
Akutsu, T., Ando, M., Arai, K., et al. 2019, Nature Astronomy, 3, 35, doi: 10.1038/s41550-018-0658-y
Allen, B. 2005, PhRvD, 71, 062001, doi: 10.1103/PhysRevD.71.062001
Allen, B., Anderson, W. G., Brady, P. R., Brown, D. A., & Creighton, J. D. E. 2012, PhRvD, 85, 122006, doi: 10.1103/PhysRevD.85.122006
Alléné, C., Aubin, F., Bentara, I., et al. 2025, Classical and Quantum Gravity, 42, 105009, doi: 10.1088/1361-6382/add234
Álvarez-López, S., Liyanage, A., Ding, J., Ng, R., & McIver, J. 2024, Classical and Quantum Gravity, 41, 085007, doi: 10.1088/1361-6382/ad2194
Andres, N., Assiduo, M., Aubin, F., et al. 2022, Classical and Quantum Gravity, 39, 055002, doi: 10.1088/1361-6382/ac482a
Arnaud, N., Driggers, J., O'Reilly, B., et al. 2026, in Journal of Physics Conference Series, Vol. 3177, Journal of Physics Conference Series (IOP), 012067, doi: 10.1088/1742-6596/3177/1/012067
Ashton, G., & Talbot, C. 2021, MNRAS, 507, 2037, doi: 10.1093/mnras/stab2236
Ashton, G., Udall, R., & Yarbrough, Z. 2022, CBC Workflow. https://git.ligo.org/cbc/projects/cbcflow
Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, ApJS, 241, 27, doi: 10.3847/1538-4365/ab06fc
Aubin, F., Brighenti, F., Chierici, R., et al. 2021, Classical and Quantum Gravity, 38, 095004, doi: 10.1088/1361-6382/abe913
Aubin, F., Bentara, I., Buskulic, D., et al. 2025, The MBTA PSD calculation for the O4a offline all-sky search, Tech. Rep. TDS-0722A. https://tds.virgo-gw.eu/?r=25024
Babak, S., Biswas, R., Brady, P. R., et al. 2013, PhRvD, 87, 024033, doi: 10.1103/PhysRevD.87.024033
Baka, T., Wright, M., Romero-Shaw, I., et al. 2025, Correcting misspecification of calibration uncertainties in gravitational-wave data analysis with efficient reweighting, Tech. Rep. LIGO-T2500295, LIGO Project. https://dcc.ligo.org/LIGO-T2500295/public
Barausse, E., Buonanno, A., Hughes, S. A., et al. 2012, Phys. Rev. D, 85, 024046, doi: 10.1103/PhysRevD.85.024046
Bernuzzi, S., Nagar, A., Dietrich, T., & Damour, T. 2015, PhRvL, 114, 161103, doi: 10.1103/PhysRevLett.114.161103
Bhaumik, S., Gayathri, V., Bartos, I., et al. 2025, Phys. Rev. D, 111, 123032, doi: 10.1103/hwr5-scp4
Biscoveanu, S., Isi, M., Varma, V., & Vitale, S. 2021, PhRvD, 104, 103018, doi: 10.1103/PhysRevD.104.103018

Blackman, J., Field, S. E., Scheel, M. A., et al. 2017a, PhRvD, 95, 104023, doi: 10.1103/PhysRevD.95.104023

—. 2017b, PhRvD, 96, 024058, doi: 10.1103/PhysRevD.96.024058

Blanchet, L. 2014, Living Reviews in Relativity, 17, 2, doi: 10.12942/lrr-2014-2

Bohé, A., Hannam, M., Husa, S., et al. 2016, PhenomPv2 - Technical Notes for LAL Implementation, Tech. Rep. LIGO-T1500602, LIGO Project. https://dcc.ligo.org/LIGO-T1500602

Bohe, A., Marsat, S., Faye, G., & Blanchet, L. 2013, Class. Quant. Grav., 30, 075017, doi: 10.1088/0264-9381/30/7/075017

Bohé, A., Shao, L., Taracchini, A., et al. 2017, PhRvD, 95, 044028, doi: 10.1103/PhysRevD.95.044028

Borchers, A., Ohme, F., Mielke, J., & Ghosh, S. 2024, Phys. Rev. D, 110, 024037, doi: 10.1103/PhysRevD.110.024037

Bose, S., Dayanga, T., Ghosh, S., & Talukder, D. 2011, Classical and Quantum Gravity, 28, 134009, doi: 10.1088/0264-9381/28/13/134009

Boyle, M., Owen, R., & Pfeiffer, H. P. 2011, PhRvD, 84, 124011, doi: 10.1103/PhysRevD.84.124011

Brown, D. A., Kumar, P., & Nitz, A. H. 2013, PhRvD, 87, 082004, doi: 10.1103/PhysRevD.87.082004

Buonanno, A., Chen, Y., & Vallisneri, M. 2003, PhRvD, 67, 104025, doi: 10.1103/PhysRevD.67.104025

Buonanno, A., & Damour, T. 1999, PhRvD, 59, 084006, doi: 10.1103/PhysRevD.59.084006

—. 2000, PhRvD, 62, 064015, doi: 10.1103/PhysRevD.62.064015

Buonanno, A., Iyer, B. R., Ochsner, E., Pan, Y., & Sathyaprakash, B. S. 2009, PhRvD, 80, 084043, doi: 10.1103/PhysRevD.80.084043

Buonanno, A., Pan, Y., Baker, J. G., et al. 2007, Phys. Rev. D, 76, 104049, doi: 10.1103/PhysRevD.76.104049

Cabero, M., Lundgren, A., Nitz, A. H., et al. 2019, Classical and Quantum Gravity, 36, 155010, doi: 10.1088/1361-6382/ab2e14

Cabourn Davies, G. S., & Harry, I. W. 2022, Classical and Quantum Gravity, 39, 215012, doi: 10.1088/1361-6382/ac8862

Cahillane, C., Betzwieser, J., Brown, D. A., et al. 2017, PhRvD, 96, 102001, doi: 10.1103/PhysRevD.96.102001

Calderón Bustillo, J., Husa, S., Sintes, A. M., & Pürrer, M. 2016, PhRvD, 93, 084019, doi: 10.1103/PhysRevD.93.084019

Calderón Bustillo, J., Laguna, P., & Shoemaker, D. 2017, PhRvD, 95, 104038, doi: 10.1103/PhysRevD.95.104038

Calderón Bustillo, J., Sanchis-Gual, N., Torres-Forné, A., & Font, J. A. 2021, Phys. Rev. Lett., 126, 201101, doi: 10.1103/PhysRevLett.126.201101

Cannon, K., Hanna, C., & Keppel, D. 2012a, PhRvD, 85, 081504, doi: 10.1103/PhysRevD.85.081504

—. 2013, PhRvD, 88, 024025, doi: 10.1103/PhysRevD.88.024025

Cannon, K., Hanna, C., & Peoples, J. 2015, arXiv e-prints, arXiv:1504.04632, doi: 10.48550/arXiv.1504.04632

Cannon, K., Cariou, R., Chapman, A., et al. 2012b, ApJ, 748, 136, doi: 10.1088/0004-637X/748/2/136

Cannon, K., Caudill, S., Chan, C., et al. 2021, SoftwareX, 14, 100680, doi: 10.1016/j.softx.2021.100680

Capano, C., Pan, Y., & Buonanno, A. 2014, Phys. Rev. D, 89, 102003, doi: 10.1103/PhysRevD.89.102003

Chandra, K., Villa-Ortega, V., Dent, T., et al. 2021, PhRvD, 104, 042004, doi: 10.1103/PhysRevD.104.042004

Chatziioannou, K. 2020, Gen. Rel. Grav., 52, 109, doi: 10.1007/s10714-020-02754-3

Chatziioannou, K., Cornish, N., Klein, A., & Yunes, N. 2015, ApJL, 798, L17, doi: 10.1088/2041-8205/798/1/L17

Chatziioannou, K., Cornish, N., Wijngaarden, M., & Littenberg, T. B. 2021, PhRvD, 103, 044013, doi: 10.1103/PhysRevD.103.044013

Chatziioannou, K., Haster, C.-J., Littenberg, T. B., et al. 2019, PhRvD, 100, 104004, doi: 10.1103/PhysRevD.100.104004

Chatziioannou, K., Klein, A., Yunes, N., & Cornish, N. 2013, Phys. Rev. D, 88, 063011, doi: 10.1103/PhysRevD.88.063011

Chatziioannou, K., Klein, A., Yunes, N., & Cornish, N. 2017, PhRvD, 95, 104004, doi: 10.1103/PhysRevD.95.104004

Chia, H. S., Edwards, T. D. P., Wadekar, D., et al. 2024, Phys. Rev. D, 110, 063007, doi: 10.1103/PhysRevD.110.063007

Chu, Q., Kovalam, M., Wen, L., et al. 2022, PhRvD, 105, 024023, doi: 10.1103/PhysRevD.105.024023

Colleoni, M., Ramis Vidal, F. A., García-Quirós, C., Akçay, S., & Bera, S. 2025a, PhRvD, 111, 104019, doi: 10.1103/PhysRevD.111.104019

Colleoni, M., Ramis Vidal, F. A., Johnson-McDaniel, N. K., et al. 2025b, PhRvD, 111, 064025, doi: 10.1103/PhysRevD.111.064025

Cornish, N. J., & Littenberg, T. B. 2015, Classical and Quantum Gravity, 32, 135012, doi: 10.1088/0264-9381/32/13/135012

Cornish, N. J., Littenberg, T. B., Bécsy, B., et al. 2021, PhRvD, 103, 044006, doi: 10.1103/PhysRevD.103.044006

Cotesta, R., Buonanno, A., Bohé, A., et al. 2018, PhRvD, 98, 084028, doi: 10.1103/PhysRevD.98.084028

Cotesta, R., Marsat, S., & Pürrer, M. 2020, PhRvD, 101, 124040, doi: 10.1103/PhysRevD.101.124040

Creighton, J. 2019, Certain Identities in FGMC, Tech. Rep. LIGO-T1700029-v2, LIGO. https://dcc.ligo.org/T1700029/public

Cutler, C., & Flanagan, É. E. 1994, PhRvD, 49, 2658, doi: 10.1103/PhysRevD.49.2658

Dal Canton, T., & Harry, I. W. 2017, arXiv e-prints, arXiv:1705.01845, doi: 10.48550/arXiv.1705.01845

Dal Canton, T., Lundgren, A. P., & Nielsen, A. B. 2015, Phys. Rev. D, 91, 062010, doi: 10.1103/PhysRevD.91.062010

Dal Canton, T., Nitz, A. H., Gadre, B., et al. 2021, ApJ, 923, 254, doi: 10.3847/1538-4357/ac2f9a

Dal Canton, T., Nitz, A. H., Lundgren, A. P., et al. 2014, PhRvD, 90, 082004, doi: 10.1103/PhysRevD.90.082004

Damour, T., Iyer, B. R., & Sathyaprakash, B. S. 2001, Phys. Rev. D, 63, 044023, doi: 10.1103/PhysRevD.63.044023

Damour, T., & Nagar, A. 2014, Phys. Rev. D, 90, 044018, doi: 10.1103/PhysRevD.90.044018

Damour, T., Nagar, A., & Villain, L. 2012, PhRvD, 85, 123007, doi: 10.1103/PhysRevD.85.123007

Dartez, L., et al. 2025, Characterization of systematic error in Advanced LIGO calibration in the fourth observing run, In preparation

Davies, G. S., Dent, T., Tápai, M., et al. 2020, PhRvD, 102, 022004, doi: 10.1103/PhysRevD.102.022004

Davis, D., Littenberg, T. B., Romero-Shaw, I. M., et al. 2022, Classical and Quantum Gravity, 39, 245013, doi: 10.1088/1361-6382/aca238

Davis, D., Yarbrough, Z., Areeda, J., et al. 2026, arXiv e-prints, arXiv:2605.16183. https://arxiv.org/abs/2605.16183

Dax, M., Green, S. R., Gair, J., et al. 2021, PhRvL, 127, 241103, doi: 10.1103/PhysRevLett.127.241103

—. 2023, PhRvL, 130, 171403, doi: 10.1103/PhysRevLett.130.171403

Dent, T. 2025, Extending the PyCBC pastro calculation to a global network, Tech. Rep. LIGO-T2100060-v3, LIGO. https://dcc.ligo.org/T2100060/public

Devine, C., Etienne, Z. B., & McWilliams, S. T. 2016, Class. Quant. Grav., 33, 125025, doi: 10.1088/0264-9381/33/12/125025

Dhurkunde, R., & Nitz, A. H. 2025, Physical Review D, 111, doi: 10.1103/physrevd.111.103018

Di Renzo, F., Fidecaro, F., Razzano, M., & Sorrentino, N. 2024, arXiv e-prints, arXiv:2401.15392, doi: 10.48550/arXiv.2401.15392

Di Renzo, F., & Virgo Collaboration. 2023, in 41st International Conference on High Energy Physics, 110

Dietrich, T., Bernuzzi, S., & Tichy, W. 2017, PhRvD, 96, 121501, doi: 10.1103/PhysRevD.96.121501

Dietrich, T., Samajdar, A., Khan, S., et al. 2019a, PhRvD, 100, 044003, doi: 10.1103/PhysRevD.100.044003

Dietrich, T., Radice, D., Bernuzzi, S., et al. 2018, Class. Quant. Grav., 35, 24LT01, doi: 10.1088/1361-6382/aaebc0

Dietrich, T., Khan, S., Dudi, R., et al. 2019b, PhRvD, 99, 024029, doi: 10.1103/PhysRevD.99.024029

Drago, M., Klimenko, S., Lazzaro, C., et al. 2021, SoftwareX, 14, 100678, doi: 10.1016/j.softx.2021.100678

Earl, D. J., & Deem, M. W. 2005, Physical Chemistry Chemical Physics (Incorporating Faraday Transactions), 7, 3910, doi: 10.1039/B509983H

Essick, R., Godwin, P., Hanna, C., Blackburn, L., & Katsavounidis, E. 2020, Machine Learning: Science and Technology, 2, 015004, doi: 10.1088/2632-2153/abab5f

Essick, R., Godwin, P., Hanna, C., Blackburn, L., & Katsavounidis, E. 2020, arXiv e-prints, arXiv:2005.12761, doi: 10.48550/arXiv.2005.12761

Essick, R., Coughlin, M. W., Zevin, M., et al. 2025, PhRvD, 112, 102001, doi: 10.1103/44x3-hv3y

Estellés, H., Buonanno, A., Enficiaud, R., Foo, C., & Pompili, L. 2026, PhRvD, 113, 044049, doi: 10.1103/pjbd-pjxn

Estellés, H., Colleoni, M., García-Quirós, C., et al. 2022a, PhRvD, 105, 084040, doi: 10.1103/PhysRevD.105.084040

Estellés, H., Husa, S., Colleoni, M., et al. 2022b, PhRvD, 105, 084039, doi: 10.1103/PhysRevD.105.084039

Estellés, H., Ramos-Buades, A., Husa, S., et al. 2021, PhRvD, 103, 124060, doi: 10.1103/PhysRevD.103.124060

Ewing, B., Huxford, R., Singh, D., et al. 2024, PhRvD, 109, 042008, doi: 10.1103/PhysRevD.109.042008

Farr, W. 2014, Marginalisation of the time and phase parameters in CBC parameter estimation, Tech. Rep. DCC-T1400460, LIGO. https://dcc.ligo.org/LIGO-T1400460/public

Farr, W., Farr, B., & Littenberg, T. 2014, Modelling calibration errors in CBC waveforms, Tech. Rep. DCC-T1400682, LIGO. https://dcc.ligo.org/LIGO-T1400682/public

Farr, W. M., Gair, J. R., Mandel, I., & Cutler, C. 2015, PhRvD, 91, 023005, doi: 10.1103/PhysRevD.91.023005

Favata, M., Kim, C., Arun, K. G., Kim, J., & Lee, H. W. 2022, PhRvD, 105, 023003, doi: 10.1103/PhysRevD.105.023003

Finn, L. S. 1992, PhRvD, 46, 5236, doi: 10.1103/PhysRevD.46.5236

Flanagan, É. É., & Hinderer, T. 2008, PhRvD, 77, 021502, doi: 10.1103/PhysRevD.77.021502

Fong, H. K. Y. 2018, PhD thesis, Toronto U.

Foucart, F., Deaton, M. B., Duez, M. D., et al. 2014, Phys. Rev. D, 90, 024026, doi: 10.1103/PhysRevD.90.024026

Gamba, R., Akçay, S., Bernuzzi, S., & Williams, J. 2022, Phys. Rev. D, 106, 024020, doi: 10.1103/PhysRevD.106.024020

Gamba, R., Bernuzzi, S., & Nagar, A. 2021, Phys. Rev. D, 104, 084058, doi: 10.1103/PhysRevD.104.084058

Gamba, R., Breschi, M., Carullo, G., et al. 2023a, Nature Astron., 7, 11, doi: 10.1038/s41550-022-01813-w

Gamba, R., et al. 2023b, Analytically improved and numerical-relativity informed effective-one-body model for coalescing binary neutron stars. https://arxiv.org/abs/2307.15125

García-Quirós, C., Colleoni, M., Husa, S., et al. 2020, PhRvD, 102, 064002, doi: 10.1103/PhysRevD.102.064002

García-Quirós, C., Husa, S., Mateu-Lucena, M., & Borchers, A. 2021, Classical and Quantum Gravity, 38, 015006, doi: 10.1088/1361-6382/abc36e

Gayathri, V., Healy, J., Lange, J., et al. 2022, Nature Astron., 6, 344, doi: 10.1038/s41550-021-01568-w

Gerosa, D., Fumagalli, G., Mould, M., et al. 2023, Phys. Rev. D, 108, 024042, doi: 10.1103/PhysRevD.108.024042

Gerosa, D., Kesden, M., Sperhake, U., Berti, E., & O'Shaughnessy, R. 2015, PhRvD, 92, 064016, doi: 10.1103/PhysRevD.92.064016

Ghonge, S., Chatziioannou, K., Clark, J. A., et al. 2020, PhRvD, 102, 064056, doi: 10.1103/PhysRevD.102.064056

Ghonge, S., Brandt, J., Sullivan, J. M., et al. 2024, Phys. Rev. D, 110, 122002, doi: 10.1103/PhysRevD.110.122002

Ghosh, S., Kolitsidou, P., & Hannam, M. 2024, Phys. Rev. D, 109, 024061, doi: 10.1103/PhysRevD.109.024061

Glanzer, J., Banagiri, S., Coughlin, S. B., et al. 2023, Classical and Quantum Gravity, 40, 065004, doi: 10.1088/1361-6382/acb633

Godwin, P. 2020, PhD thesis, Pennsylvania State University

Godwin, P., Essick, R., Hanna, C., et al. 2020, arXiv e-prints, arXiv:2010.15282, doi: 10.48550/arXiv.2010.15282

Goggans, P. M., & Chi, Y. 2004, in American Institute of Physics Conference Series, Vol. 707, Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. G. J. Erickson & Y. Zhai (AIP), 59–66, doi: 10.1063/1.1751356

Gonzalez, A., Gamba, R., Breschi, M., et al. 2023a, Phys. Rev. D, 107, 084026, doi: 10.1103/PhysRevD.107.084026

Gonzalez, A., et al. 2023b, Class. Quant. Grav., 40, 085011, doi: 10.1088/1361-6382/acc231

Graff, P. B., Buonanno, A., & Sathyaprakash, B. S. 2015, PhRvD, 92, 022002, doi: 10.1103/PhysRevD.92.022002

Green, S. R., Simpson, C., & Gair, J. 2020, PhRvD, 102, 104057, doi: 10.1103/PhysRevD.102.104057

Guo, X., Chu, Q., Chung, S. K., et al. 2018, Computer Physics Communications, 231, 62, doi: 10.1016/j.cpc.2018.05.002

Gupta, T., & Cornish, N. J. 2024, PhRvD, 109, 064040, doi: 10.1103/PhysRevD.109.064040

Gupte, N., et al. 2024, Evidence for eccentricity in the population of binary black holes observed by LIGO-Virgo-KAGRA. https://arxiv.org/abs/2404.14286

Hamilton, E., London, L., Thompson, J. E., et al. 2021, PhRvD, 104, 124027, doi: 10.1103/PhysRevD.104.124027

Hamilton, E., Colleoni, M., Thompson, J. E., et al. 2026, PhRvD, 113, 084055, doi: 10.1103/kxsf-23rr

Hanna, C., Caudill, S., Messick, C., et al. 2020, PhRvD, 101, 022003, doi: 10.1103/PhysRevD.101.022003

Hanna, C., et al. 2022, Phys. Rev. D, 106, 084033, doi: 10.1103/PhysRevD.106.084033

Hanna, C., Kennington, J., Sakon, S., et al. 2023, PhRvD, 108, 042003, doi: 10.1103/PhysRevD.108.042003

Hannam, M., Schmidt, P., Bohé, A., et al. 2014, PhRvL, 113, 151101, doi: 10.1103/PhysRevLett.113.151101

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2

Harris, F. 1978, Proceedings of the IEEE, 66, 51, doi: 10.1109/PROC.1978.10837

Harry, I., & Hinderer, T. 2018, Class. Quant. Grav., 35, 145010, doi: 10.1088/1361-6382/aac7e3

Harry, I., Privitera, S., Bohé, A., & Buonanno, A. 2016, Phys. Rev. D, 94, 024012, doi: 10.1103/PhysRevD.94.024012

Harry, I. W., Allen, B., & Sathyaprakash, B. S. 2009, PhRvD, 80, 104014, doi: 10.1103/PhysRevD.80.104014

Harry, I. W., & Fairhurst, S. 2011, Physical Review D, 83, doi: 10.1103/physrevd.83.084002

Hastings, W. K. 1970, Biometrika, 57, 97, doi: 10.1093/biomet/57.1.97

Healy, J., & Lousto, C. O. 2017, PhRvD, 95, 024037, doi: 10.1103/PhysRevD.95.024037

Helmling-Cornell, A. F., Nguyen, P., Schofield, R. M. S., & Frey, R. 2024, Classical and Quantum Gravity, 41, 145003, doi: 10.1088/1361-6382/ad5139

Henry, Q. 2023, Phys. Rev. D, 107, 044057, doi: 10.1103/PhysRevD.107.044057

Henry, Q., Faye, G., & Blanchet, L. 2020, Phys. Rev. D, 102, 044033, doi: 10.1103/PhysRevD.102.044033

Hinderer, T. 2008, Astrophys. J., 677, 1216, doi: 10.1086/533487

Hinderer, T., Taracchini, A., Foucart, F., et al. 2016, PhRvL, 116, 181101, doi: 10.1103/PhysRevLett.116.181101

Hofmann, F., Barausse, E., & Rezzolla, L. 2016, ApJL, 825, L19, doi: 10.3847/2041-8205/825/2/L19

Hooper, S., Chung, S. K., Luan, J., et al. 2012, PhRvD, 86, 024012, doi: 10.1103/PhysRevD.86.024012

Hourihane, S., Chatziioannou, K., Wijngaarden, M., et al. 2022, PhRvD, 106, 042006, doi: 10.1103/PhysRevD.106.042006

Hoy, C., & Raymond, V. 2021, SoftwareX, 15, 100765, doi: 10.1016/j.softx.2021.100765

Hunter, J. D. 2007, Computing in Science and Engineering, 9, 90, doi: 10.1109/MCSE.2007.55

Husa, S., Khan, S., Hannam, M., et al. 2016, PhRvD, 93, 044006, doi: 10.1103/PhysRevD.93.044006

Islam, T., Field, S. E., & Khanna, G. 2023, Phys. Rev. D, 108, 064048, doi: 10.1103/PhysRevD.108.064048

Isogai, T., LIGO Scientific Collaboration, & Virgo Collaboration. 2010, in Journal of Physics Conference Series, Vol. 243, Journal of Physics Conference Series (IOP), 012005, doi: 10.1088/1742-6596/243/1/012005

Isoyama, S., Sturani, R., & Nakano, H. 2020, doi: 10.1007/978-981-15-4702-7_31-1

Jan, A., Tsao, B.-J., O'Shaughnessy, R., Shoemaker, D., & Laguna, P. 2025. https://arxiv.org/abs/2508.12460

Jiménez-Forteza, X., Keitel, D., Husa, S., et al. 2017, PhRvD, 95, 064024, doi: 10.1103/PhysRevD.95.064024

Johnson-McDaniel, N. K., Ghosh, A., Ghonge, S., et al. 2022a, PhRvD, 105, 044020, doi: 10.1103/PhysRevD.105.044020

Johnson-McDaniel, N. K., Gupta, A., Ajith, P., et al. 2016, Determining the final spin of a binary black hole system including in-plane spins: Method and checks of accuracy, Tech. Rep. DCC-T1600168, LIGO. https://dcc.ligo.org/LIGO-T1600168/public

Johnson-McDaniel, N. K., Kulkarni, S., & Gupta, A. 2022b, PhRvD, 106, 023001, doi: 10.1103/PhysRevD.106.023001

Joshi, P., Tsukada, L., & Hanna, C. 2023, PhRvD, 108, 084032, doi: 10.1103/PhysRevD.108.084032

Joshi, P., Tsukada, L., Hanna, C., et al. 2025a, arXiv e-prints, arXiv:2506.06497, doi: 10.48550/arXiv.2506.06497

Joshi, P., Niu, W., Hanna, C., et al. 2025b, arXiv e-prints, arXiv:2505.23959, doi: 10.48550/arXiv.2505.23959

Kacanja, K., Soni, K., & Nitz, A. H. 2025. https://arxiv.org/abs/2508.00179

Kapadia, S. J., Caudill, S., Creighton, J. D. E., et al. 2020, Classical and Quantum Gravity, 37, 045007, doi: 10.1088/1361-6382/ab5f2d

Kawaguchi, K., Kyutoku, K., Nakano, H., et al. 2015, Phys. Rev. D, 92, 024014, doi: 10.1103/PhysRevD.92.024014

Keitel, D., Forteza, X. J., Husa, S., et al. 2017, PhRvD, 96, 024006, doi: 10.1103/PhysRevD.96.024006

Khalil, M., Buonanno, A., Estelles, H., et al. 2023, Phys. Rev. D, 108, 124036, doi: 10.1103/PhysRevD.108.124036

Khan, S., Husa, S., Hannam, M., et al. 2016, PhRvD, 93, 044007, doi: 10.1103/PhysRevD.93.044007

Khan, S., Ohme, F., Chatziioannou, K., & Hannam, M. 2020, PhRvD, 101, 024056, doi: 10.1103/PhysRevD.101.024056

Kiuchi, K., Kawaguchi, K., Kyutoku, K., et al. 2017, Phys. Rev. D, 96, 084060, doi: 10.1103/PhysRevD.96.084060

Klein, A., Cornish, N., & Yunes, N. 2013, Phys. Rev. D, 88, 124015, doi: 10.1103/PhysRevD.88.124015

—. 2014, Phys. Rev. D, 90, 124029, doi: 10.1103/PhysRevD.90.124029

Klimenko, S. 2022, arXiv e-prints, arXiv:2201.01096, doi: 10.48550/arXiv.2201.01096

Klimenko, S., Mohanty, S., Rakhmanov, M., & Mitselmakher, G. 2005, PhRvD, 72, 122002, doi: 10.1103/PhysRevD.72.122002

Klimenko, S., Yakushin, I., Mercer, A., & Mitselmakher, G. 2008, Classical and Quantum Gravity, 25, 114029, doi: 10.1088/0264-9381/25/11/114029

Klimenko, S., Vedovato, G., Drago, M., et al. 2016, PhRvD, 93, 042004, doi: 10.1103/PhysRevD.93.042004

Kolitsidou, P., Thompson, J. E., & Hannam, M. 2024, Impact of anti-symmetric contributions to signal multipoles in the measurement of black-hole spins. https://arxiv.org/abs/2402.00813

Kovalam, M., Kaium Patwary, M. A., Sreekumar, A. K., et al. 2022, ApJL, 927, L9, doi: 10.3847/2041-8213/ac5687

Krishnendu, N. V., & Ohme, F. 2022, PhRvD, 105, 064012, doi: 10.1103/PhysRevD.105.064012

Krolak, A., & Schutz, B. F. 1987, General Relativity and Gravitation, 19, 1163, doi: 10.1007/BF00759095

Kumar, P., & Dent, T. 2024, PhRvD, 110, 043036, doi: 10.1103/PhysRevD.110.043036

Kumar, S., Nitz, A. H., & Forteza, X. J. 2025, The Astrophysical Journal, 982, 67, doi: 10.3847/1538-4357/adb973

Kyutoku, K., Okawa, H., Shibata, M., & Taniguchi, K. 2011, Phys. Rev. D, 84, 064018, doi: 10.1103/PhysRevD.84.064018

Lackey, B. D., Pürrer, M., Taracchini, A., & Marsat, S. 2019, PhRvD, 100, 024002, doi: 10.1103/PhysRevD.100.024002

Lang, R. N., & Hughes, S. A. 2006, Phys. Rev. D, 74, 122001, doi: 10.1103/PhysRevD.75.089902

Lange, J., O'Shaughnessy, R., Boyle, M., et al. 2017, PhRvD, 96, 104041, doi: 10.1103/PhysRevD.96.104041

Lenon, A. K., Nitz, A. H., & Brown, D. A. 2020, MNRAS, 497, 1966, doi: 10.1093/mnras/staa2120

LIGO–Virgo–KAGRA Collaboration. 2018, LVK Algorithm Library - LALSuite, Free software (GPL), doi: 10.7935/GT1W-FZ16

—. 2025a, GWTC-4.0: Candidate Data Release, Zenodo, doi: 10.5281/zenodo.17014083

—. 2025b, LIGO/Virgo/KAGRA Public Alerts User Guide. https://emfollow.docs.ligo.org/userguide

—. 2025c, GWTC-4.0: Data Quality Products for Transient Gravitational Wave Searches, Zenodo, doi: 10.5281/zenodo.16856919

—. 2025d, GWTC-4.0: Parameter Estimation Data Release, Zenodo, doi: 10.5281/zenodo.17014085

—. 2026, GWTC-5.0: Glitch Modelling for Events, Zenodo, doi: XX.YYY/zenodo.ZZZZZZZZ

LIGO Scientific Collaboration, Virgo Collaboration, & KAGRA collaboration. 2022, Noise curves used for Simulations in the update of the Observing Scenarios Paper, Tech. Rep. LIGO-T2000012. https://dcc.ligo.org/LIGO-T2000012/public

Lin, J. 1991, IEEE Transactions on Information Theory, 37, 145, doi: 10.1109/18.61115

Littenberg, T. B., & Cornish, N. J. 2009, PhRvD, 80, 063007, doi: 10.1103/PhysRevD.80.063007

—. 2015, PhRvD, 91, 084034, doi: 10.1103/PhysRevD.91.084034

Liu, Y., Du, Z., Chung, S. K., et al. 2012, Classical and Quantum Gravity, 29, 235018, doi: 10.1088/0264-9381/29/23/235018

London, L., Khan, S., Fauchon-Jones, E., et al. 2018, PhRvL, 120, 161102, doi: 10.1103/PhysRevLett.120.161102

Lower, M. E., Thrane, E., Lasky, P. D., & Smith, R. 2018, PhRvD, 98, 083028, doi: 10.1103/PhysRevD.98.083028

Luan, J., Hooper, S., Wen, L., & Chen, Y. 2012, PhRvD, 85, 102002, doi: 10.1103/PhysRevD.85.102002

Maggiore, M. 2007, Gravitational Waves: Volume 1: Theory and Experiments, doi: 10.1093/acprof:oso/9780198570745.001.0001

Marsat, S., Blanchet, L., Bohe, A., & Faye, G. 2013, Gravitational waves from spinning compact object binaries: New post-Newtonian results. https://arxiv.org/abs/1312.5375

Martel, K., & Poisson, E. 1999, PhRvD, 60, 124008, doi: 10.1103/PhysRevD.60.124008

Matas, A., Dietrich, T., Buonanno, A., et al. 2020, PhRvD, 102, 043023, doi: 10.1103/PhysRevD.102.043023

Mehta, A. K., Wadekar, D., Roulet, J., et al. 2025. https://arxiv.org/abs/2501.17939

Messick, C., Blackburn, K., Brady, P., et al. 2017, PhRvD, 95, 042001, doi: 10.1103/PhysRevD.95.042001

Mihaylov, D. P., Ossokine, S., Buonanno, A., et al. 2023, pySEOBNR: a software package for the next generation of effective-one-body multipolar waveform models. https://arxiv.org/abs/2303.18203

Mihaylov, D. P., Ossokine, S., Buonanno, A., & Ghosh, A. 2021, Phys. Rev. D, 104, 124087, doi: 10.1103/PhysRevD.104.124087

Mills, C., & Fairhurst, S. 2021, PhRvD, 103, 024042, doi: 10.1103/PhysRevD.103.024042

Mishra, T., Bhaumik, S., Gayathri, V., et al. 2025, PhRvD, 111, 023054, doi: 10.1103/PhysRevD.111.023054

Mishra, T., O'Brien, B., Gayathri, V., et al. 2021, Phys. Rev. D, 104, 023014, doi: 10.1103/PhysRevD.104.023014

Mishra, T., O'Brien, B., Szczepańczyk, M., et al. 2022, PhRvD, 105, 083018, doi: 10.1103/PhysRevD.105.083018

Moe, B., Brady, P., Stephens, B., et al. 2014, GraceDB: A Gravitational Wave Candidate Event Database, Tech. Rep. LIGO-T1400365-v5, LIGO Scientific Collaboration. https://dcc.ligo.org/LIGO-T1400365/public

Morras, G., Pratten, G., & Schmidt, P. 2025a, Phys. Rev. D, 111, 084052, doi: 10.1103/PhysRevD.111.084052

—. 2025b. https://arxiv.org/abs/2503.15393

Mould, M., & Gerosa, D. 2022, Phys. Rev. D, 105, 024076, doi: 10.1103/PhysRevD.105.024076

Mozzon, S., Ashton, G., Nuttall, L. K., & Williamson, A. R. 2022, Phys. Rev. D, 106, 043504, doi: 10.1103/PhysRevD.106.043504

Mozzon, S., Nuttall, L. K., Lundgren, A., et al. 2020, Classical and Quantum Gravity, 37, 215014, doi: 10.1088/1361-6382/abac6c

Mukherjee, D., Caudill, S., Magee, R., et al. 2021, PhRvD, 103, 084047, doi: 10.1103/PhysRevD.103.084047

Nagar, A., & Rettegno, P. 2019, Phys. Rev. D, 99, 021501, doi: 10.1103/PhysRevD.99.021501

Nagar, A., & Shah, A. 2016, Phys. Rev. D, 94, 104017, doi: 10.1103/PhysRevD.94.104017

Nagar, A., et al. 2018, Phys. Rev. D, 98, 104052, doi: 10.1103/PhysRevD.98.104052

Necula, V., Klimenko, S., & Mitselmakher, G. 2012, in Journal of Physics Conference Series, Vol. 363, Journal of Physics Conference Series (IOP), 012032, doi: 10.1088/1742-6596/363/1/012032

Nitz, A. H. 2018, Classical and Quantum Gravity, 35, 035016, doi: 10.1088/1361-6382/aaa13d

Nitz, A. H., Capano, C., Nielsen, A. B., et al. 2019, ApJ, 872, 195, doi: 10.3847/1538-4357/ab0108

Nitz, A. H., Dal Canton, T., Davis, D., & Reyes, S. 2018, PhRvD, 98, 024050, doi: 10.1103/PhysRevD.98.024050

Nitz, A. H., Dent, T., Dal Canton, T., Fairhurst, S., & Brown, D. A. 2017, ApJ, 849, 118, doi: 10.3847/1538-4357/aa8f50

Nitz, A. H., Dent, T., Davies, G. S., et al. 2020, ApJ, 891, 123, doi: 10.3847/1538-4357/ab733f

Nuttall, L. K. 2018, Philosophical Transactions of the Royal Society of London Series A, 376, 20170286, doi: 10.1098/rsta.2017.0286

O'Shaughnessy, R., Healy, J., London, L., Meeks, Z., & Shoemaker, D. 2012, Phys. Rev. D, 85, 084003, doi: 10.1103/PhysRevD.85.084003

O'Shea, E., & Kumar, P. 2023, PhRvD, 108, 104018, doi: 10.1103/PhysRevD.108.104018

Ossokine, S., Buonanno, A., Marsat, S., et al. 2020, PhRvD, 102, 044055, doi: 10.1103/PhysRevD.102.044055

Owen, B. J. 1996, PhRvD, 53, 6749, doi: 10.1103/PhysRevD.53.6749

Pan, Y., Buonanno, A., Taracchini, A., et al. 2014, PhRvD, 89, 084006, doi: 10.1103/PhysRevD.89.084006

Pankow, C., Brady, P., Ochsner, E., & O'Shaughnessy, R. 2015, PhRvD, 92, 023002, doi: 10.1103/PhysRevD.92.023002

Pankow, C., Chatziioannou, K., Chase, E. A., et al. 2018, PhRvD, 98, 084016, doi: 10.1103/PhysRevD.98.084016

Pannarale, F., Berti, E., Kyutoku, K., Lackey, B. D., & Shibata, M. 2015, Phys. Rev. D, 92, 084050, doi: 10.1103/PhysRevD.92.084050

Pannarale, F., Berti, E., Kyutoku, K., & Shibata, M. 2013, Phys. Rev. D, 88, 084011, doi: 10.1103/PhysRevD.88.084011

Pannarale, F., Rezzolla, L., Ohme, F., & Read, J. S. 2011, PhRvD, 84, 104017, doi: 10.1103/PhysRevD.84.104017

Payne, E., Talbot, C., Lasky, P. D., Thrane, E., & Kissel, J. S. 2020, PhRvD, 102, 122004, doi: 10.1103/PhysRevD.102.122004

Peters, P. C. 1964, Physical Review, 136, 1224, doi: 10.1103/PhysRev.136.B1224

Peters, P. C., & Mathews, J. 1963, Physical Review, 131, 435, doi: 10.1103/PhysRev.131.435

Phukon, K. S., Schmidt, P., & Pratten, G. 2025, Phys. Rev. D, 111, 043040, doi: 10.1103/PhysRevD.111.043040

Planas, M. d. L., Husa, S., Ramos-Buades, A., & Valencia, J. 2025a. https://arxiv.org/abs/2506.01760

Planas, M. d. L., Ramos-Buades, A., García-Quirós, C., et al. 2025b, Eccentric or circular? A reanalysis of binary black hole gravitational wave events for orbital eccentricity signatures. https://arxiv.org/abs/2504.15833

Poisson, E. 1998, PhRvD, 57, 5287, doi: 10.1103/PhysRevD.57.5287

Pompili, L., Buonanno, A., Estellés, H., et al. 2023, PhRvD, 108, 124035, doi: 10.1103/PhysRevD.108.124035

Powell, J. 2018, Classical and Quantum Gravity, 35, 155017, doi: 10.1088/1361-6382/aacf18

Pratten, G., Husa, S., García-Quirós, C., et al. 2020a, PhRvD, 102, 064001, doi: 10.1103/PhysRevD.102.064001

Pratten, G., Schmidt, P., Buscicchio, R., & Thomas, L. M. 2020b, Physical Review Research, 2, 043096, doi: 10.1103/PhysRevResearch.2.043096

Pratten, G., García-Quirós, C., Colleoni, M., et al. 2021, PhRvD, 103, 104056, doi: 10.1103/PhysRevD.103.104056

Privitera, S., Mohapatra, S. R. P., Ajith, P., et al. 2014, PhRvD, 89, 024003, doi: 10.1103/PhysRevD.89.024003

Ramos-Buades, A., Buonanno, A., Estellés, H., et al. 2023, PhRvD, 108, 124037, doi: 10.1103/PhysRevD.108.124037

Ramos-Buades, A., Husa, S., Pratten, G., et al. 2020, Phys. Rev. D, 101, 083015, doi: 10.1103/PhysRevD.101.083015

Ramos-Buades, A., Schmidt, P., Pratten, G., & Husa, S. 2020a, PhRvD, 101, 103014, doi: 10.1103/PhysRevD.101.103014

Ramos-Buades, A., Tiwari, S., Haney, M., & Husa, S. 2020b, PhRvD, 102, 043005, doi: 10.1103/PhysRevD.102.043005

Ray, A., Niu, W., Sakon, S., et al. 2023, arXiv e-prints, arXiv:2306.07190, doi: 10.48550/arXiv.2306.07190

Rodriguez, C. L., Zevin, M., Pankow, C., Kalogera, V., & Rasio, F. A. 2016, ApJL, 832, L2, doi: 10.3847/2041-8205/832/1/L2

Romero-Shaw, I., Lasky, P. D., Thrane, E., & Calderón Bustillo, J. 2020a, ApJL, 903, L5, doi: 10.3847/2041-8213/abbe26

Romero-Shaw, I., Stegmann, J., Tagawa, H., et al. 2025, Phys. Rev. D, 112, 063052, doi: 10.1103/jj7m-x66y

Romero-Shaw, I. M., Gerosa, D., & Loutrel, N. 2023, Mon. Not. Roy. Astron. Soc., 519, 5352, doi: 10.1093/mnras/stad031

Romero-Shaw, I. M., Talbot, C., Biscoveanu, S., et al. 2020b, MNRAS, 499, 3295, doi: 10.1093/mnras/staa2850

Roy, S., Sengupta, A. S., & Ajith, P. 2019, PhRvD, 99, 024048, doi: 10.1103/PhysRevD.99.024048

Roy, S., Sengupta, A. S., & Thakor, N. 2017, PhRvD, 95, 104045, doi: 10.1103/PhysRevD.95.104045

Sachdev, S., Caudill, S., Fong, H., et al. 2019, arXiv e-prints, arXiv:1901.08580, doi: 10.48550/arXiv.1901.08580

Sakon, S., Tsukada, L., Fong, H., et al. 2024, PhRvD, 109, 044066, doi: 10.1103/PhysRevD.109.044066

Salemi, F., Milotti, E., Prodi, G. A., et al. 2019, PhRvD, 100, 042003, doi: 10.1103/PhysRevD.100.042003

Salpeter, E. E. 1955, Astrophys. J., 121, 161, doi: 10.1086/145971

Samajdar, A., & Dietrich, T. 2018, Phys. Rev. D, 98, 124030, doi: 10.1103/PhysRevD.98.124030

Samsing, J., & Ramirez-Ruiz, E. 2017, Astrophys. J. Lett., 840, L14, doi: 10.3847/2041-8213/aa6f0b

Santamaría, L., Ohme, F., Ajith, P., et al. 2010, PhRvD, 82, 064016, doi: 10.1103/PhysRevD.82.064016

Schmidt, P., Hannam, M., & Husa, S. 2012, PhRvD, 86, 104063, doi: 10.1103/PhysRevD.86.104063

Schmidt, P., Hannam, M., Husa, S., & Ajith, P. 2011, PhRvD, 84, 024046, doi: 10.1103/PhysRevD.84.024046

Schmidt, P., Ohme, F., & Hannam, M. 2015, PhRvD, 91, 024043, doi: 10.1103/PhysRevD.91.024043

Singer, L. P., & Price, L. R. 2016, PhRvD, 93, 024013, doi: 10.1103/PhysRevD.93.024013

Singer, L. P., Chen, H.-Y., Holz, D. E., et al. 2016, ApJL, 829, L15, doi: 10.3847/2041-8205/829/1/L15

Singer, L. P., et al. 2016, Astrophys. J. Suppl., 226, 10, doi: 10.3847/0067-0049/226/1/10

Skilling, J. 2006, Bayesian Analysis, 1, 833, doi: 10.1214/06-BA127

Smith, J. R., Abbott, T., Hirose, E., et al. 2011, Classical and Quantum Gravity, 28, 235005, doi: 10.1088/0264-9381/28/23/235005

Smith, R. J. E., Ashton, G., Vajpeyi, A., & Talbot, C. 2020, MNRAS, 498, 4492, doi: 10.1093/mnras/staa2483

Soni, S., Glanzer, J., Effler, A., et al. 2024, Class. Quant. Grav., 41, 135015, doi: 10.1088/1361-6382/ad494a

Soni, S., Austin, C., Effler, A., et al. 2021, Classical and Quantum Gravity, 38, 025016, doi: 10.1088/1361-6382/abc906

Soni, S., Berger, B. K., Davis, D., et al. 2025, Classical and Quantum Gravity, 42, 085016, doi: 10.1088/1361-6382/adc4b6

Speagle, J. S. 2020, MNRAS, 493, 3132, doi: 10.1093/mnras/staa278

Steinhoff, J., Hinderer, T., Buonanno, A., & Taracchini, A. 2016, PhRvD, 94, 104028, doi: 10.1103/PhysRevD.94.104028

Steinle, N., & Kesden, M. 2022, Phys. Rev. D, 106, 063028, doi: 10.1103/PhysRevD.106.063028

Stevenson, S., Berry, C. P. L., & Mandel, I. 2017, MNRAS, 471, 2801, doi: 10.1093/mnras/stx1764

Stovall, K., Freire, P. C. C., Chatterjee, S., et al. 2018, ApJL, 854, L22, doi: 10.3847/2041-8213/aaad06

Sturani, R. 2015, Note on the derivation of the angular momentum and spin precessing equations in SpinTaylor codes, Tech. Rep. DCC-T1500554, LIGO. https://dcc.ligo.org/LIGO-T1500554/public

Sun, L., Goetz, E., Kissel, J. S., et al. 2020, Classical and Quantum Gravity, 37, 225008, doi: 10.1088/1361-6382/abb14e

—. 2021, arXiv e-prints, arXiv:2107.00129, doi: 10.48550/arXiv.2107.00129

Talbot, C., & Thrane, E. 2017, PhRvD, 96, 023012, doi: 10.1103/PhysRevD.96.023012

Talbot, C., Biscoveanu, S., Zimmerman, A., et al. 2025, Classical and Quantum Gravity, 42, 235023, doi: 10.1088/1361-6382/ae1ac7

Tantau, T. 2023, The TikZ and PGF Packages. https://github.com/pgf-tikz/pgf

Taracchini, A., Buonanno, A., Khanna, G., & Hughes, S. A. 2014a, Phys. Rev. D, 90, 084025, doi: 10.1103/PhysRevD.90.084025

Taracchini, A., et al. 2014b, Phys. Rev. D, 89, 061502, doi: 10.1103/PhysRevD.89.061502

Thompson, J. E., Fauchon-Jones, E., Khan, S., et al. 2020, PhRvD, 101, 124059, doi: 10.1103/PhysRevD.101.124059

Thompson, J. E., Hamilton, E., London, L., et al. 2024, PhRvD, 109, 063012, doi: 10.1103/PhysRevD.109.063012

Thrane, E., & Talbot, C. 2019, PASA, 36, e010, doi: 10.1017/pasa.2019.2

Tiwari, V., Klimenko, S., Necula, V., & Mitselmakher, G. 2016, Classical and Quantum Gravity, 33, 01LT01, doi: 10.1088/0264-9381/33/1/01LT01

Tsukada, L., Joshi, P., Adhicary, S., et al. 2023, PhRvD, 108, 043004, doi: 10.1103/PhysRevD.108.043004

Tucker, A., & Will, C. M. 2021, PhRvD, 104, 104023, doi: 10.1103/PhysRevD.104.104023

Usman, S. A., Nitz, A. H., Harry, I. W., et al. 2016, Classical and Quantum Gravity, 33, 215004, doi: 10.1088/0264-9381/33/21/215004

Vajente, G. 2022, PhRvD, 105, 102005, doi: 10.1103/PhysRevD.105.102005

van de Meent, M., Buonanno, A., Mihaylov, D. P., et al. 2023, Phys. Rev. D, 108, 124038, doi: 10.1103/PhysRevD.108.124038

van der Sluys, M., Raymond, V., Mandel, I., et al. 2008, Classical and Quantum Gravity, 25, 184011, doi: 10.1088/0264-9381/25/18/184011

Varma, V., Field, S. E., Scheel, M. A., et al. 2019, Physical Review Research, 1, 033015, doi: 10.1103/PhysRevResearch.1.033015

Varma, V., Gerosa, D., Stein, L. C., Hébert, F., & Zhang, H. 2019, Phys. Rev. Lett., 122, 011101, doi: 10.1103/PhysRevLett.122.011101

Varma, V., Isi, M., & Biscoveanu, S. 2020, Phys. Rev. Lett., 124, 101104, doi: 10.1103/PhysRevLett.124.101104

Vazsonyi, L., & Davis, D. 2023, Classical and Quantum Gravity, 40, 035008, doi: 10.1088/1361-6382/acafd2

Vecchio, A. 2004, Phys. Rev. D, 70, 042001, doi: 10.1103/PhysRevD.70.042001

Veitch, J., Pürrer, M., & Mandel, I. 2015a, PhRvL, 115, 141101, doi: 10.1103/PhysRevLett.115.141101

Veitch, J., Raymond, V., Farr, B., et al. 2015b, PhRvD, 91, 042003, doi: 10.1103/PhysRevD.91.042003

Villa-Ortega, V., Dent, T., & Barroso, A. C. 2022, Mon. Not. Roy. Astron. Soc., 515, 5718, doi: 10.1093/mnras/stac2120

Vinciguerra, S., Veitch, J., & Mandel, I. 2017, Class. Quant. Grav., 34, 115006, doi: 10.1088/1361-6382/aa6d44

Vines, J., Flanagan, É. É., & Hinderer, T. 2011, PhRvD, 83, 084051, doi: 10.1103/PhysRevD.83.084051

Vitale, S., Lynch, R., Veitch, J., Raymond, V., & Sturani, R. 2014, PhRvL, 112, 251101, doi: 10.1103/PhysRevLett.112.251101

Wadekar, D., Venumadhav, T., Roulet, J., et al. 2024, PhRvD, 110, 044063, doi: 10.1103/PhysRevD.110.044063

Warburton, N., Pound, A., Wardell, B., Miller, J., & Durkan, L. 2021, Phys. Rev. Lett., 127, 151102, doi: 10.1103/PhysRevLett.127.151102

Wen, L. 2003, Astrophys. J., 598, 419, doi: 10.1086/378794

—. 2008, International Journal of Modern Physics D, 17, 1095–1104, doi: 10.1142/s0218271808012723

Wette, K. 2020, SoftwareX, 12, 100634, doi: 10.1016/j.softx.2020.100634

Williams, D., Veitch, J., Chiofalo, M., et al. 2023, The Journal of Open Source Software, 8, 4170, doi: 10.21105/joss.04170

Wysocki, D., O'Shaughnessy, R., Lange, J., & Fang, Y.-L. L. 2019, PhRvD, 99, 084026, doi: 10.1103/PhysRevD.99.084026

Yu, H., Roulet, J., Venumadhav, T., Zackay, B., & Zaldarriaga, M. 2023, Phys. Rev. D, 108, 064059, doi: 10.1103/PhysRevD.108.064059

Zackay, B., Venumadhav, T., Roulet, J., Dai, L., & Zaldarriaga, M. 2021, PhRvD, 104, 063034, doi: 10.1103/PhysRevD.104.063034

Zappa, F., Bernuzzi, S., Pannarale, F., Mapelli, M., & Giacobbo, N. 2019, PhRvL, 123, 041102, doi: 10.1103/PhysRevLett.123.041102

Zhu, X., Thrane, E., Oslowski, S., Levin, Y., & Lasky, P. D. 2018, Phys. Rev. D, 98, 043002, doi: 10.1103/PhysRevD.98.043002