MASTER IN ASTROPHYSICS AND SPACE SCIENCE

Erasmus Mundus Master

in Astrophysics and Space Science

Master Thesis

Bayesian model classification for burst gravitational-wave signals

BASILIC: rapid, reproducible burst science

Supervisors

Doct. Marie Anne Bizouard

CNRS, Université Côte d'Azur, Observatoire de la Côte d'Azur

Prof. Viviana Fafone

Università degli studi di Roma, Tor Vergata

Author

Iuliu Daniel Cuceu

Academic Year 2024/2025

TOR VERGATA UNIVERSITÀ DEGLI STUDI DI ROMA

UNIVERSITY OF BELGRADE

University of Bremen

UNIVERSITÉ CÔTE D'AZUR

Co-funded by the European Union

This Master thesis is submitted in partial fulfillment of the requirements for the degree Master Physique Fondamental et Applications, parcours MASS as part of a multiple degree awarded in the framework of the Erasmus Mundus Joint Master in Astrophysics and Space Science −− MASS jointly delivered by a Consortium of four Universities: Tor Vergata University of Rome, University of Belgrade, University of Bremen, and Université Cote d'Azur, regulated by the MASS Consortium Agreement and funded by the EU under the call ERASMUS-EDU-2021-PEX-EMJM-MOB.

Abstract

Gravitational-wave astronomy now routinely detects compact binary coalescences. Yet, short-duration burst signals which are brief, information-rich transients from a variety of sources, remain challenging to classify. Beyond detector sensitivity, the bottleneck is practical, with robust Bayesian parameter estimation and model comparison requiring bespoke code, careful bookkeeping, and access to high-throughput computing. This thesis presents BASILIC, an end-to-end pipeline that turns a short configuration file into a complete Bayesian analysis and standardised report. BASILIC builds on the `bilby` inference library, ships with a curated registry of burst and compact-binary waveform models, and integrates natively with HTCondor for scalable single-event and injection-campaign studies. We illustrate two use cases to showcase the pipeline's scientific reach. First, we analyse the 23 November 2023 event, GW231123, under multiple hypotheses. BASILIC recovers binary-black-hole parameters consistent with the LVK results, and goes further, by showing that power-law and cosmic-string burst morphologies, while favoured over pure noise, are strongly disfavoured relative to a binary-black-hole origin. Second, a large injection study quantifies parameter recovery and model discrimination across signal-to-noise ratio and morphology, highlighting regions of genuine confusion, e.g., very high-mass binary-black-holes versus low-index power-law bursts, where multiple hypotheses can achieve comparable evidence. By standardising inputs, orchestration, and outputs, BASILIC lowers the barrier to rapid, hypothesis-aware gravitational wave burst science and provides a practical foundation for systematic model comparison as detector sensitivity and event rates increase.

Contents

Introduction

The direct detection of gravitational waves (GWs) has transformed our ability to observe compact, energetic phenomena in the Universe. While all detections so far are well described by compact binary coalescences (CBCs), a broad class of short-duration "burst" signals remains undiscovered. Bursts can arise from diverse mechanisms, such as relativistic features on cosmic strings, core-collapse supernovae, gravitational-wave memory, neutron star oscillation or the late stages of very massive black-hole mergers. The common theme is the presence of only a handful of in-band cycles. This brevity complicates both detection and interpretation: different physical scenarios can generate similar time-frequency morphologies, and low signal-to-noise ratios exacerbate ambiguity.

From a practical standpoint, robust Bayesian inference, estimating parameters and, crucially, model comparison, exists within the LIGO [1], Virgo [2] and KAGRA [3] Collaboration (LVK), but it is not yet trivial to deploy at scale for bursts. A typical single-event study can involve: selecting and conditioning data; choosing a power spectral density (PSD); wiring likelihoods, priors, and waveform generators; managing computing jobs; and finally producing plots and tables that are consistent across hypotheses. Injection campaigns multiply this complexity, demanding reproducible noise realisations, careful resource allocation, and aggregation of many runs into statistically meaningful summaries. These steps impose a high entry cost: expertise in coding and high-throughput computing can become a limiting factor in doing the science.

This thesis addresses that bottleneck by introducing Basilic—BAyesian tranSIent modeLIng and Classification—an end-to-end pipeline designed to make model-aware GW burst analyses routine. The design is guided by the principles of no user coding, leveraging pre-existing software, scalability through construction and reproducibility.

Scientifically, Basilic ships with a curated set of waveforms spanning both physically motivated and phenomenological models. On the CBC side, it supports commonly used waveforms for high-mass binary black holes (BBH). For burst sources, it includes frequency-domain power laws and cosmic-string cusp/kink families, memory, simplified transients (e.g., sine-Gaussian, damped sinusoids), and phenomenological supernova-inspired templates. This breadth enables hypothesis testing be-

yond a single favoured model, which is essential for short, loud events with few in-band cycles. In short, BASILIC reframes GW burst inference from a bespoke, code-heavy exercise into a reproducible pipeline: write one file, submit, and read a coherent, hypothesis-aware report. As event rates increase and morphology diversity widens, such tooling will be key to turning candidate detections into quantitative, multi-model science.

The remainder of the thesis is organised as follows. Chapter 1 reviews essential background: the nature of GW burst signals and the Bayesian tools used to analyse them. Chapter 2 presents BASILIC's design, user experience, and current feature set, emphasizing the minimal inputs required and the standardised outputs produced. Chapter 3 illustrates two use cases. First, a multi-hypothesis re-analysis of GW231123 [4] shows that binary-black-hole models are decisively preferred over burst alternatives, despite those alternatives outperforming pure noise. This particular analysis resulted in a manuscript currently under review by Physical Review Letters [5]. Second, an injection campaign quantifies recoverability and model confusion across signal morphologies and signal-to-noise ratios, identifying regimes where even strong signals can be ambiguous between families (e.g., very high-mass BBHs versus low-index power-law bursts).

1 Burst gravitational wave signals

The discovery of gravitational waves (GWs) has ushered in a new era of astronomy-one that listens, rather than looks, at the Universe. Rooted in Einstein's general theory of relativity, gravitational waves are ripples in spacetime generated by accelerating masses. They propagate at the speed of light and interact only weakly with matter, enabling them to carry pristine information from some of the most extreme and distant events in the cosmos.

Of particular interest, is a subset of gravitational wave signals, called bursts or transients, typically understood as signals shorter than at most a couple of seconds, with many astrophysically rich examples being present in the detector sensitivity range for a dozen milliseconds or so. The challenge with such signals comes from both a theoretical perspective, different sources with rich physical variety can conspire to produce similar signals, but also from a data analysis perspective, since a thorough statistical treatment of this potential variety, in the presence of large noise levels, is anything but trivial.

In order to tackle these challenges systematically, in this chapter, we first introduce the fundamental aspects of gravitational waves, starting from their mathematical formulation in the linearised Einstein equations, their propagation in spacetime, and their physical effects. For a more complete description of gravitational wave theory, please see e.g. [6]. Second, we will approach the data analysis methods used and challenges faced when decoding different gravitational wave signals, from matched filtering and Bayesian statistics, to numerical methods like Nested Sampling [7, 8]. The aim is to provide a brief yet descriptive, introduction to the gravitational wave theory and data analysis methods that will be used in future chapters.

1.1 Listening with gravitational waves

This section provides a brief overview of the theoretical framework that underpins gravitational wave physics. We present a review of the physical mechanisms responsible for gravitational wave generation. We cover a variety of astrophysical and cosmological sources, all of which have parametrised models that are part of Basilic, and thus can be used for hypothesis testing, as explained in Section 2. The main references for this section are [6, 9–11].

1.1.1 The nature of gravitational waves

Einstein's field equations govern the curvature of spacetime induced by energy and momentum. In vacuum, they take the form:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 0, \tag{1}$$

where $R_{\mu\nu}$, R is the Ricci tensor and Ricci scalar respectively. In the weak-field approximation, we consider small perturbations $h_{\mu\nu}$ around a flat Minkowski background:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}, \quad |h_{\mu\nu}| \ll 1. \tag{2}$$

Here, $\eta_{\mu\nu}$ is the Minkowski metric, and $h_{\mu\nu}$ represents small deviations. Substituting this into the field equations and keeping terms only to first order in $h_{\mu\nu}$, we obtain the linearised Einstein equations:

$$\Box\bar{h}_{\mu\nu} = 0, \quad \bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h, \tag{3}$$

where $\Box = \eta^{\alpha\beta}\partial_\alpha\partial_\beta$ is the d'Alembertian operator, and $h = h^\alpha_\alpha$ is the trace of the perturbation. This equation describes the propagation of a wave-like perturbation in spacetime, analogous to the electromagnetic wave equation in vacuum.

Plane wave solution and transverse nature

Solving this wave equation in vacuum yields a familiar plane wave solution:

$$h_{\mu\nu}(t, \mathbf{x}) = A_{\mu\nu}e^{ik_\alpha x^\alpha}. \tag{4}$$

Since gravity is a gauge theory, we can impose the harmonic (de Donder) gauge condition, $\partial^\nu\bar{h}_{\mu\nu} = 0$. Further refining the gauge conditions leads to the Transverse-Traceless (TT) gauge. The TT gauge simplifies the analysis of gravitational waves by reducing the ten degrees of freedom after imposing the harmonic gauge, down to only two. The conditions for the TT gauge are:

- The perturbation is purely spatial: $h_{0\mu} = 0$.
- The trace of the perturbation vanishes: $h^i_i = 0$.
- The wave is transverse: $\partial^j h_{ij} = 0$.

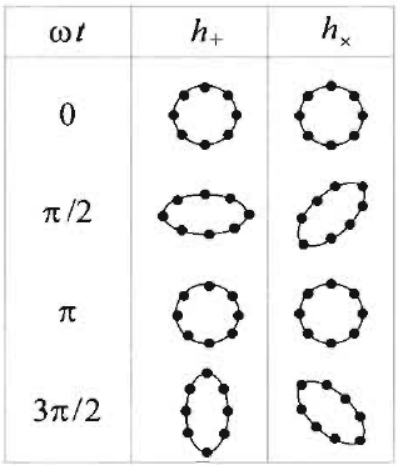

Figure 1: A ring of test particles being deformed due to plus or cross polarized passing GW, adapted from [6].

Applying these conditions for a wave propagating along z, the only non-zero components are:

$$h_{ij}^{TT}(t,z) = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix} e^{i\omega(t-z)}. \tag{5}$$

Here, h_+ and $h_\times$ are the two polarization states of gravitational waves. This is the "value-of-interest" in source modelling problems, thus a parametrization of h_+ and $h_\times$ for some GW-emitter, usually given in the TT gauge, is sufficient to completely describe the signal morphology.

Physical effects and energy transport

Gravitational waves induce tidal forces on test masses. Consider a ring of test particles in the xy-plane. The effect of an h_+ polarized wave is to stretch and squeeze the ring along the x- and y-axes alternately. An $h_\times$ wave causes a similar effect, but rotated by 45°, see Fig. 1. The energy carried by GWs is given by the Isaacson stress-energy tensor:

$$T_{\mu\nu}^{GW} = \frac{1}{32\pi G}\langle \partial_\mu h_{\alpha\beta} \partial_\nu h^{\alpha\beta} \rangle. \tag{6}$$

This expression, averaged over several wavelengths, describes how GWs transport energy and momentum across spacetime.

1.1.2 Detection principles

We have introduced gravitational waves as metric perturbations in the TT gauge. However, an Earth-based detector does not move freely in spacetime; rather, it is anchored to the Earth and influenced by local forces. This requires transitioning to a "detector-adapted frame", where we analyse how GWs influence test masses and interferometer arms.

The ideal situation would be a drag-free satellite, free-falling with respect to the total gravitational field, given by both the Earth and the GWs present, such as using Fermi normal coordinates. For an Earth-bound detector however, it would have an acceleration $a = -g$ with respect to a local inertial frame, and it would rotate relative to local gyroscopes. In this frame, the metric takes the form:

$$\begin{aligned} ds^2 \simeq & -c^2dt^2\Big[1+\frac{2}{c^2}\boldsymbol{a}\cdot\boldsymbol{x}+\frac{1}{c^4}(\boldsymbol{a}\cdot\boldsymbol{x})^2-\frac{1}{c^2}(\boldsymbol{\Omega}\times\boldsymbol{x})^2+R_{oioj}x^ix^j\Big] \\ & +2cdt\,dx^i\Big[\frac{1}{c}\,\epsilon_{ijk}\,\Omega^j x^k-\frac{2}{3}\,R_{0jik}\,x^jx^k\Big] \\ & +dx^idx^j\Big[\delta_{ij}-\frac{1}{3}\,R_{ikjl}\,x^kx^l\Big] \end{aligned} \tag{7}$$

One can immediately read off the "physical" meaning of the term $2\boldsymbol{a}\cdot\boldsymbol{x}/c^2$ as the inertial acceleration. Similarly, the term $(\boldsymbol{a}\cdot\boldsymbol{x}/c^2)^2$ is a gravitational redshift, while the term $(\boldsymbol{\Omega}\times\boldsymbol{x}/c)^2$ is a Lorentz time dilation due to the angular velocity of the Earth-bound laboratory. Lastly, $(1/c)\epsilon_{ijk}\Omega^j x^k$ is known as the Sagnac effect, describing the phase difference in an interferometer that is rotating. The terms proportional to the Riemann tensor describe the effects of the slowly varying gravitational background, and the effect of GWs. Expanding this in terms of r/L, where $r^2 = x^ix^i$ and L is a typical variation scale of the metric, to zeroth order we would be describing Newtonian physics. This allows for an interpretation of the various terms as "Newtonian forces", which is more intuitive.

Interaction with a test mass

Consider a freely-falling test mass in a GW background. The geodesic deviation equation governs how nearby masses are displaced under the influence of curvature:

$$\frac{D^2\xi^i}{D\tau^2} = -R^i{}_{kjl}\xi^j\frac{dx^k}{d\tau}\frac{dx^l}{d\tau}, \tag{8}$$

where ξ^i is the displacement of a test mass, $R^i{}_{kjl}$ is the Riemann tensor, and we have the covariant derivative w.r.t an affine parameter, $\frac{D^2}{D\tau^2}$. Since the Riemann tensor is

invariant in linearised theory, it is simpler to compute the non-zero components in the TT gauge, even if we are studying the effects in the proper detector frame. For a GW propagating along the z-direction, the relevant (non-zero) Riemann tensor components are:

$$R_{0x0x} = -R_{0y0y} = -\frac{1}{2}\frac{\partial^2 h_+}{\partial t^2}, \quad R_{0x0y} = R_{0y0x} = -\frac{1}{2}\frac{\partial^2 h_\times}{\partial t^2}. \tag{9}$$

Thus, the displacement of a test mass follows:

$$\ddot{\xi}^i = \frac{1}{2}\ddot{h}^{\mathrm{TT}}_{ij}\xi^j, \tag{10}$$

As stated previously, this shows that GWs induce "tidal accelerations", causing a Newtonian-like effect on test masses. The dimensionless strain h represents the fractional change in separation between test masses due to a GW. Typical astrophysical sources, such as binary black hole mergers, produce strains on the order of $h \sim 10^{-21}$, for a BBH at 1 Gpc. Such small perturbations require extremely sensitive detection techniques, leading to the development of laser interferometers.

Basic Michelson interferometers

Interferometric GW detectors, such as LIGO [1] and Virgo [2], measure changes in the length of perpendicular arms caused by passing gravitational waves. They operate by using laser interferometry to monitor the phase shift induced by spacetime perturbations. A standard Michelson interferometer, see Fig. 2, consists of: a laser producing a coherent light source; a beam splitter that divides the laser beam into two perpendicular arms; mirrors at the ends of the arms, reflecting the beams back to the beam splitter; a photodetector, which measures the interference pattern from the recombined beams.

Gravitational waves alter the relative arm lengths, shifting the interference pattern. In the absence of a GW, the interference is ideally destructive, resulting in minimal light at the photodetector. A passing GW changes the relative optical path lengths, leading to a measurable change in the detected light intensity. To improve sensitivity, interferometers include Fabry-Perot cavities in each arm. These consist of partially reflective mirrors that allow light to bounce multiple times, effectively increasing the optical path length. If a GW-induced displacement ΔL is present, the effective phase shift is amplified by the number of reflections.

Stable operation requires cleaning and recycling the light. Input and output mode cleaners transmit only the fundamental Gaussian spatial mode and suppress

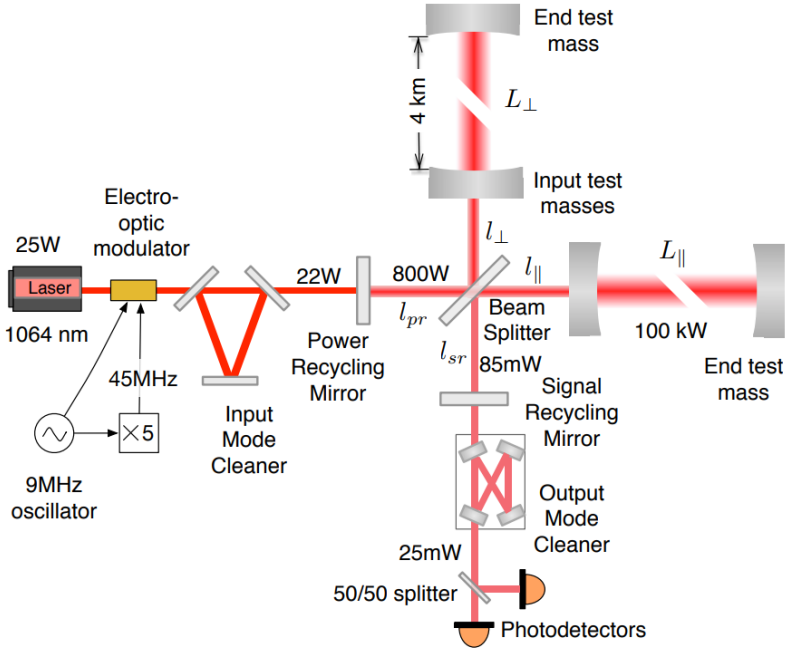

Figure 2: Schematic of the layout of an Advanced LIGO detector. The Michelson interferometer is enhanced by two 4-km resonant arm cavities. The input mode cleaner is a triangular Fabry-Perot cavity, and the power-recycling mirror resonates the laser's light to increase the power further. The signal recycling mirror on the other hand is used to broaden the response of the detector and lastly the output mode cleaner rejects unwanted spatial and frequency components of the light before detection. Illustration taken from ref. [12].

beam jitter and higher-order modes. A power-recycling mirror forms a resonant cavity with the interferometer to raise the circulating power, increasing sensitivity. A signal-recycling mirror at the antisymmetric port shapes the frequency response and allows a tunable trade between bandwidth and peak sensitivity.

Sensitivity gains come with noise trade-offs. At high frequencies the dominant limit is photon shot noise. At low frequencies radiation pressure from the same photons impart stochastic forces on the test masses. These two quantum effects set a standard quantum limit that depends on circulating power and mirror mass. Thermal noise from test-mass coatings and suspensions contributes broadly across the band and is reduced with low-loss materials and careful suspension design. Feedback control needed to hold the interferometer on resonance introduces sensing and actuation noise that must be filtered and isolated from the science channel.

Below about 20 Hz environmental effects dominate. Seismic motion couples through the suspension chain, while Newtonian noise arises from local density fluc-

tuations in the ground that directly perturb the test masses. Multi-stage isolation, quiet sites, and underground facilities mitigate these effects. A representative noise budget is shown in Fig. 3, where measured amplitude spectral densities are compared with a modelled sum of quantum, thermal, residual-gas, seismic, alignment and length-control, beam-jitter, laser, photodetector dark, actuator, and damping contributions.

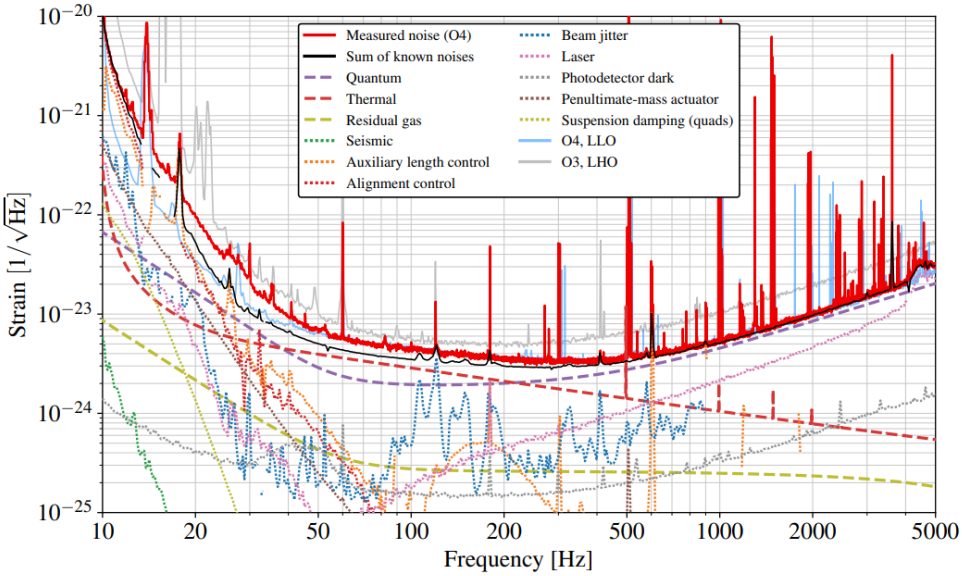

Figure 3: Amplitude spectral density of the LIGO Hanford detector with a modelled noise budget. Individual contributions include quantum (shot noise and radiation pressure), thermal (test-mass coatings and suspensions), residual gas, seismic, auxiliary length and alignment control, beam jitter, laser, photodetector dark, actuator, and suspension damping. The model sum is shown together with the measured noise in O4, and for reference the O4 Livingston and O3 Hanford curves; adapted from Fig. 5(a) of [11].

Recorded data: Time series format

After all these layers of optical enhancement and noise mitigation, the final output of the detector is a calibrated time-domain signal known as the "strain", representing the relative length change between interferometer arms. The data is stored as a time series:

$$s(t) = h(t) + n(t), \tag{11}$$

where $s(t)$ is the detector output, $h(t)$ is the gravitational wave signal, and $n(t)$ represents the cumulative noise from all sources. This data format enables various

analysis strategies. The calibrated strain is passed through pipelines involving, matched filtering, Fourier analysis and Bayesian inference, among others, which we will look at in detail in Sec. 1.2.

1.1.3 Gravitational wave sources

The potential astrophysical phenomenology responsible for the production of GW is varied, and sometimes not well understood. In this subsection, we will go over some of the common culprits that can produce GW, which with current-generation ground-based detectors we classify as burst signals. The goal is to introduce the waveforms, specific equations describing the plus and cross polarization of a travelling GW, of the models we will use in subsequent chapters. We begin with the best understood, and observationally confirmed source compact binary coalescence, for which high stellar-mass ($\gtrsim 10\,M_\odot$) black hole collisions are the main candidate of burst signals, followed by cosmic strings, memory effects, and lastly, core-collapse supernovae.

Compact binary coalescence

CBCs are the most well-modelled and observationally established sources of gravitational waves. They encompass the inspiral, merger, and ringdown of two compact objects, typically black holes (BHs), neutron stars (NSs), or a combination thereof. The modelling of these events has been one of the triumphs of general relativity in both its analytical and numerical formulations. The theoretical understanding of CBC waveforms hinges on the interplay between various approximation techniques: the post-Newtonian (PN) expansion for the inspiral, numerical relativity (NR) for the merger, and black hole perturbation theory for the ringdown. Please refer to, e.g. [13], for a more thorough introduction of the theory and phenomenology we briefly outline below.

Gravitational waves are emitted by systems with a time-varying mass quadrupole moment. This contrasts with electromagnetism, where a time-varying dipole moment suffices to produce radiation. The leading order gravitational radiation is given by the celebrated quadrupole formula:

$$h_{ij}(t, \boldsymbol{x}) = \frac{2G}{c^4 r}\ddot{Q}_{ij}(t - r/c), \tag{12}$$

where Q_{ij}^{TT} is the TT-gauge projection of the source's mass quadrupole tensor and r

is the distance to the source. For a binary system in circular orbit, the quadrupole moment oscillates with twice the orbital frequency, $\omega_{gw} = 2\omega_{\text{orbit}}$, resulting in a waveform that increases in both frequency and amplitude, the characteristic "chirp". This is the main effect controlling the duration of the signal inside detector, and why we can classify a subset of BBH CBCs as burst signals. To illustrate this, let us consider two point masses m_1 and m_2 in a circular orbit separated by a distance R, with the reduced mass $\mu = m_1 m_2/(m_1 + m_2)$, and total mass, M. The quadrupole tensor for such a system reads:

$$Q_{ij} = \mu \left(x_i x_j - \frac{1}{3}\delta_{ij} R^2 \right), \tag{13}$$

where $\mathbf{x}$ is the separation vector between the two masses. From the quadrupole emission, one can derive the rate of energy loss due to gravitational radiation:

$$\frac{dE}{dt} = -\frac{32}{5}\frac{c^5}{G}\left(\frac{G\mathcal{M}_c \omega_{gw}}{2c^3}\right)^{10/3}, \tag{14}$$

where we have introduced the chirp mass $\mathcal{M}_c = \mu^{3/5} M^{2/5}$. This energy loss back-reacts on the orbit, leading to an inspiral, the slow decrease in orbital separation and increase in frequency. The system's evolution is governed by, $\frac{d\omega}{dt} \propto \omega^{11/3}$, indicating a rapidly accelerating chirp as the binary components spiral inwards. Therefore, a system "starts" inspiraling in frequency band far below current generation detectors, and as the masses get closer and frequency increases, the system becomes "audible" to us. The higher the system mass, the lower the coalescence frequency, and the shorter the signal appears in the interferometer. This effect is a major player in the injection campaigns presented in Chapter 3. This behaviour sets the stage for the post-Newtonian modelling of the waveform, a method which systematically improves the quadrupole approximation using a velocity expansion $v/c \ll 1$.

The PN expansion is a perturbative method for solving the Einstein field equations in the regime of slow-motion and weak fields. It is particularly well-suited for modelling the inspiral phase of binaries, where the velocities are still significantly below the speed of light and the gravitational fields are not yet extreme. In this regime, the metric is expanded as:

$$g_{\mu\nu} = \eta_{\mu\nu} + h^{(2)}_{\mu\nu} + h^{(4)}_{\mu\nu} + h^{(6)}_{\mu\nu} + \ldots, \tag{15}$$

where $h^{(n)}_{\mu\nu} \sim \mathcal{O}(v/c)^n$. The PN expansion is commonly carried out up to 3.5PN order, depending on the quantity studied. The conservative dynamics (i.e., neglecting

radiation reaction) are obtained from a PN-expanded Hamiltonian or Lagrangian. For example, the 2PN-order Lagrangian includes terms up to $\mathcal{O}(v^4/c^4)$ and captures periastron precession, higher-order gravitational redshift effects, and spin-orbit coupling when spins are included. Radiation reaction terms enter at 2.5PN and higher orders, allowing the modelling of the gradual inspiral due to energy loss via gravitational radiation.

In the waveform domain, the PN formalism allows construction of the GW strain $h(t)$ as a sum of harmonics of the orbital frequency. The dominant $\ell = 2, m = 2$ mode is often used, but higher multipoles become important in asymmetric mass systems. One of the most used approximants derived from PN theory is the TaylorF2 waveform, constructed in the frequency domain via a stationary phase approximation. This model incorporates phase evolution to high PN orders but neglects amplitude corrections beyond leading order. Its simplicity allows for fast template generation, making it ideal for matched-filter searches, although its accuracy degrades near the merger.

As the binary evolves toward merger, the two compact objects draw nearer and their velocities approach relativistic speeds. In this strong-field, high-velocity regime, the post-Newtonian (PN) approximation breaks down. Specifically, the PN series becomes divergent or poorly convergent as $v/c \sim 0.3 - 0.5$, and non-perturbative gravitational dynamics, such as horizon formation, tidal deformation, and relativistic recoil, become essential. Accurate modelling in this regime requires solving the full nonlinear Einstein equations numerically: this is the domain of Numerical Relativity (NR) [14].

The field of NR saw a major breakthrough in 2005, when Pretorius [15], and independently the groups of Campanelli [16] and Baker [17], developed stable methods for evolving binary black hole spacetimes through merger using generalized harmonic and moving puncture formulations. These simulations solve Einstein's equations on a 3+1 spacetime decomposition, often employing the BSSN (Baumgarte-Shapiro-Shibata-Nakamura) formalism, with high-order finite differencing schemes and adaptive mesh refinement.

NR simulations are computationally expensive, often requiring weeks of supercomputing time for a single waveform, and are therefore impractical for large-scale template banks. Nevertheless, they are essential for calibrating semi-analytical models, verifying approximants, and understanding the physics of the merger phase.

Once the two black holes merge, the resulting object settles into a perturbed Kerr black hole. The dynamics of this phase are governed by black hole perturbation theory. The dominant response of a perturbed black hole is a set of damped sinusoids known as *quasi-normal modes* (QNMs), each characterized by a complex frequency $\omega_{\ell mn} = \omega_R + i\omega_I$, where the real part determines the oscillation frequency and the imaginary part determines the damping time. These modes depend only on the mass and spin of the final black hole, a manifestation of the no-hair theorem. The gravitational waveform in the ringdown phase can thus modelled as:

$$h(t) = \sum_{\ell,m,n} A_{\ell mn} e^{-i\omega_{\ell mn} t}, \tag{16}$$

where $A_{\ell mn}$ are excitation amplitudes determined by the merger. The ringdown is essential not only for understanding the late-time waveform, but also as a tool for testing general relativity. Deviations in the QNM spectrum could indicate the presence of exotic compact objects or modifications to gravity.

To enable gravitational wave data analysis, one needs waveform models that span the full coalescence: inspiral, merger, and ringdown. These models, or *approximants*, must balance physical accuracy with computational speed. They are built by stitching together the PN inspiral, NR merger, and perturbative ringdown into unified prescriptions for the GW strain $h(t)$ or its Fourier transform $\tilde{h}(f)$. There are three principal families of approximants:

- **Phenomenological Models (Phenom)**: These include the IMRPhenomD and IMRPhenomPv2 families. They construct analytic fits to PN and NR waveforms in the frequency domain. For example, IMRPhenomD models non-precessing binaries by smoothly connecting inspiral, merger, and ringdown phases via analytical templates calibrated to NR data. Precession is handled approximately via a rotation of the non-precessing frame (PhenomPv2). These models are fast and well-suited for large parameter-spaces [18–20].

- **Effective One Body (EOB) Models**: These models, such as SEOBNRv4 and SEOBNRv5, reframe the two-body problem into an effective one-body system with a deformed Schwarzschild or Kerr background. The EOB Hamiltonian is constructed to reproduce PN results at low frequencies and is extended to strong fields with NR-informed corrections. EOB models often work

in the time domain and require the integration of ordinary differential equations, making them slower but often more accurate in waveform phasing and amplitude. Ringdown is modelled by attaching QNMs to the inspiral-plunge waveform [21–23].

- **Hybrid Waveforms**: In some cases, waveforms are built by stitching together PN and NR waveforms directly, typically in the time domain. These are not analytic models per se, but they serve as benchmarks or training data for surrogate models.

Each approximant comes with limitations. PN-based models break down before merger; NR coverage is sparse in some regions of parameter space (e.g., high mass ratio, high spin); and analytic fits must be recalibrated when new physical effects (like eccentricity or higher harmonics) are included. Moreover, the waveform models are typically defined either in the time domain (e.g., EOB) or frequency domain (e.g., TaylorF2, Phenom), with implications for their computational speed and compatibility with analysis pipelines. The choice of approximant affects detection efficiency, parameter estimation, and model selection, making it a central element of any GW data analysis framework. The CBC waveform models discussed here form the backbone of many burst analyses as well, since high-mass BBH CBCs are known-events, especially when the true source may resemble a merger but with incomplete or exotic information.

Gravitational waves from cosmic strings

Cosmic strings are hypothetical $1+1$ topological defects that may have formed during symmetry-breaking phase transitions in the early Universe, predicted in theories beyond the Standard Model of particle physics [24]. Unlike the everyday intuition of "strings" in mechanics, cosmic strings are not made of matter per se; instead, they are configurations of spacetime itself, possessing mass per unit length and immense tension characterised by a dimensionless parameter $G\mu$, where G is Newton's constant and μ the string tension.

Although their existence has not yet been confirmed, cosmic strings are compelling candidates for gravitational wave bursts accessible by ground-based interferometers [25, 26]. They are expected to evolve into a scaling network of long strings and loops. It is from these loops that gravitational waves are radiated. The loops

form via inter-commutation events, when two strings cross and exchange segments, and subsequently decay via GW emission. Because the strings can move at relativistic speeds and carry large amounts of energy, certain features on the loops can emit powerful bursts of gravitational radiation.

The most distinctive features responsible for gravitational wave bursts are "cusps" and "kinks". A cusp is a point on the string that momentarily travels at the speed of light, resulting in a strongly beamed burst of GWs emitted along a narrow cone. A kink, on the other hand, is a discontinuity in the string's tangent vector that propagates at the speed of light along the string, generating a less directional, but still distinctive GW signal. The resulting bursts from these phenomena are brief, broadband, and linearly polarized, making them natural targets for burst searches with interferometers.

From an observational standpoint, cosmic-string GW bursts are especially exciting for several reasons. First, their brief and broadband nature makes them suitable for unmodelled or minimally modelled burst pipelines. Second, cosmic strings are a target for both current ground-based detectors and future space-based missions such as LISA [27], which is expected to be particularly sensitive to the low-frequency bursts emitted by large cosmic-string loops in the early Universe. Furthermore, detection of such signals would provide not only evidence for cosmic strings, but also direct insights into the energy scales of early Universe physics.

Given their importance, it is essential to construct practical waveform models for cosmic-string GW bursts. These signals can be well approximated in the frequency domain by power-law forms with sharp cut-off at high frequency. We summarise below the models we have used, taken from [28]:

For cusps, the gravitational wave burst has a characteristic frequency-domain strain:

$$\tilde{h}_{\rm cusp}(f) = A_c \Theta(f - f_{\rm low})\Theta(f_{\rm high} - f) e^{-2\pi i f t_A} f^{-4/3}, \tag{17}$$

where A_c is the amplitude parameter, t_A is the arrival time of the burst, f_h is the high-frequency cut-off determined by the opening angle of the beamed emission, and Θ is the Heaviside step function. The low frequency cut-off $f_{\rm low}$ is usually fixed at 20 Hz to take into account the noise increase at low frequency. Similarly, for kinks, the GW strain is modelled as:

$$\tilde{h}_{\rm kink}(f) = A_k \Theta(f - f_{\rm low})\Theta(f_{\rm high} - f) e^{-2\pi i f t_A} f^{-5/3}, \tag{18}$$

where A_k is the amplitude and the other parameters are as in the cusp case. The steeper spectral index $-5/3$ reflects the smoother discontinuity of a kink compared to a cusp, leading to reduced high-frequency content.

Recognising the variety and degeneracy of possible string-like signals, a general power-law model is useful:

$$\tilde{h}_p(f) = A_p \Theta(f - f_{\text{low}})\Theta(f_{\text{high}} - f) e^{-2\pi i f t_A} f^{-\alpha}, \tag{19}$$

where α is a free spectral index and A_p the amplitude, allowing this model to interpolate between cusp-like and kink-like behaviour, and even resemble other broadband burst sources such as memory effects. This generic form allows for minimal modelling while capturing key morphological features of the signal. The amplitude terms are, in the case of cusp, kink or kink-kink cosmic strings, encoding multiple physically relevant parameters, including the string tension $G\mu$, however due to mainly modulating only the amplitude rather than the phase of the signal, or frequency, they are a number of parameters whose degeneracy could not be broken easily. Thus, we write only a generic amplitude term rather than the full physically-motivated relations.

These models provide low-dimensional, phenomenologically rich templates for use in Bayesian inference pipelines. These parametrizations are capable of distinguishing between different power-law bursts using the spectral index α, and they are computationally efficient to simulate and analyse.

Gravitational wave memory effects

Gravitational wave memory refers to a permanent change in the relative separation of test masses caused by a burst of gravitational radiation. Unlike the oscillatory strains from standard GWs, which cause transient deformations that revert to zero after the wave has passed, memory effects result in a lasting displacement. This phenomenon is a striking manifestation of the nonlinear structure of general relativity and is a robust, generic prediction for a wide class of high-energy astrophysical events, particularly those involving large, unbound momentum fluxes [29, 30]. The memory effect can be broadly divided into two categories: linear memory and nonlinear memory:

- **Linear memory**: This arises from a change in the time derivative of the source's multipole moments. A canonical example is a hyperbolic scattering

event, where mass or momentum is radiated anisotropically [31]. When the incoming and outgoing states differ in their total momentum configuration, the resulting gravitational radiation causes a permanent shift in the spacetime metric. Linear memory is particularly relevant in processes like neutron star mergers, where equations of state [32] effects play a role, or supernovae where matter is ejected.

- **Nonlinear memory**: More profound is the nonlinear memory effect, arising from the stress-energy of gravitational waves themselves. Since GWs carry energy and momentum, their own flux contributes to the curvature of spacetime. The memory from this mechanism can be calculated by integrating the square of the strain tensor over time. First identified by Christodoulou in 1991 [29] and independently by Blanchet & Damour [33], this effect highlights the inherently nonlinear nature of GR, as it results in gravitational waves sourcing further gravitational waves. In both cases, the key physical insight is that gravitational wave memory encodes the cumulative energy and momentum carried away by the source. Importantly, the memory grows monotonically during the burst and plateaus after it ends, leading to a non-oscillatory, slowly varying strain.

Detecting memory poses unique challenges [34]. The signal is very low-frequency, typically below the most sensitive band of ground-based detectors. However, as noted in [28], if the burst is sufficiently energetic and compact in time, the associated memory can fall within the accessible band of detectors such as LIGO and Virgo. Furthermore, the memory has a distinctive morphology, monotonically rising and saturating, which allows it to be modelled with minimal assumptions.

To capture this slow-rise, broadband, non-oscillatory feature, the following phenomenological frequency-domain model is introduced [28]:

$$\tilde{h}_m(f) = -i\pi A_m \tau e^{-2\pi i f t_A} / \sinh\left(\pi^2 \tau f\right) \tag{20}$$

where, A_m is the memory amplitude and τ is the characteristic rise time, a function of the total mass of the astrophysical system. In the limit of $\tau \to 0$ we approach a power-law-like signal with a negative spectral index.

The simplicity of this form makes it suitable for burst data analysis pipelines. While the memory itself is non-oscillatory in time, the frequency-domain representation is nonetheless amenable to standard tools such as matched filtering and

Bayesian inference, as long as the data window includes the relevant low-frequency content. Importantly, the inclusion of such models allows for the identification of memory signals either as stand-alone burst candidates or as supplements to CBC events. In the latter case, the memory may appear as a secondary component following the merger and ringdown of a binary black hole system.

Core-collapse supernovae

Core-collapse supernovae (CCSN) are among the most energetic astrophysical events in the Universe, representing the explosive death of massive stars with initial masses greater than roughly $8\,M_\odot$. These cataclysmic explosions offer an extraordinary opportunity for multi-messenger astronomy, as they emit not only electromagnetic radiation and neutrinos, but also gravitational waves, making them potential sources for advanced ground-based detectors, although the rarity of the event has made their detection with GW elusive so far. The complexity of the physical mechanisms underlying CCSN explosions, however, makes modelling their gravitational-wave emission an extraordinarily challenging task. The main references for this part are [35–37].

The gravitational-wave signal from a CCSN is primarily emitted during the collapse of the stellar core, the bounce at nuclear densities, and the subsequent proto-neutron star (PNS) oscillations. The nature of the emitted waveform depends sensitively on the stellar structure, rotation rate, microphysics (such as the equation of state), and the post-bounce hydrodynamic instabilities, including convection and the standing accretion shock instability (SASI).

Despite the relative success in understanding the core dynamics, the resulting GW waveforms from CCSNe are stochastic and not straightforwardly amenable to matched filtering. This stochasticity is primarily due to turbulent and chaotic flows during the post-bounce evolution, especially those associated with convection and SASI. As a result, time series waveforms exhibit complex, unpredictable phase variations that render them unsuitable for template searches.

However, many of these features manifest in consistent structures within the spectrogram, which plots frequency as a function of time. Dominant emission components include high-frequency gravitational(g)- and fundamental(f)-modes associated with oscillations of the PNS surface, and low-frequency SASI modes. The key insight from recent modelling efforts, see e.g. [38], is that while the time-domain

waveforms are largely stochastic, the spectrograms reveal coherent structures that can be phenomenologically modelled.

To capture the dominant GW emission mode in a CCSN spectrogram, a new signal model called the *asymmetric chirplet* has been introduced [38]. This time-domain model accounts for the characteristic features of CCSN waveforms, including a rising frequency and asymmetric amplitude envelope. The model is defined as follows:

$$h_+(t) = A \exp\left(-b\frac{(t-t_0)^2}{\tau^2}\right) \cos\left(2\pi f(t-t_0) + 2\pi \dot{f}(t-t_0)^2 + c(t-t_0)^3\right), \quad (21)$$

$$h_\times(t) = A \exp\left(-b\frac{(t-t_0)^2}{\tau^2}\right) \sin\left(2\pi f(t-t_0) + 2\pi \dot{f}(t-t_0)^2 + c(t-t_0)^3\right), \quad (22)$$

where A is the overall amplitude; f is the peak frequency of the mode; $\dot{f}$ as the rate of frequency change (chirp); c is a cubic chirp-like parameter or curvature, allowing the frequency to plateau, as seen in simulations; $\tau = Q/(2\pi f)$, the characteristic duration, with Q the quality factor (number of cycles); and $b = 1 - a\tanh(t-t_0)$, a time-asymmetric envelope function, where a is derived from different decay rates on either side of the peak. The asymmetry parameter a is defined via decay constants α_L and α_R for the left and right sides of the chirplet, $a = \frac{\alpha_L - \alpha_R}{\alpha_L + \alpha_R}$.

This asymmetric chirplet model has been shown to match the dominant emission mode seen in numerical simulations for a range of progenitor masses and explosion scenarios. Its parameters, especially the peak frequency and duration, can be mapped onto physical properties of the PNS, such as its mass, radius, and internal structure.

To complement the chirplet model with a dedicated description of the deterministic core-bounce phase in rapidly rotating CCSNe, we adopt the three-peak analytical template proposed in [39]. The goal is to capture the short ($\lesssim 10\,\mathrm{ms}$) axisymmetric bounce morphology with the smallest number of parameters needed for inference on rotation and equation of state (EOS) proxies. The model represents the core-bounce strain as a sum of three co-spaced Gaussian peaks with amplitudes that depend on the differential rotation parameter $\beta \equiv T/|W|$ and a phenomenological shape parameter α that absorbs EOS and rotation-profile effects. Since $h_\times \approx 0$ is assumed, we model:

$$h_+(t) = h_1(\beta)\,\exp\left[-\frac{(t-\tau)^2}{2\eta^2}\right] + h_2(\beta)\,\exp\left[-\frac{(t-\tau_a)^2}{2\eta^2}\right] + h_3(\alpha,\beta)\,\exp\left[-\frac{(t-\tau_b)^2}{2\eta^2}\right], \tag{23}$$

with a common width $\eta = 0.2\,\mathrm{ms}$ and peak times $\tau_a = \tau + 0.5\,\mathrm{ms}$ and $\tau_b = \tau + 1.0\,\mathrm{ms}$, so the three extrema are equally spaced after the first arrival $\tau \in [-0.5, -0.2\,\mathrm{ms}$. The first two peak amplitudes are quadratic in β, while the third peak carries the extra EOS/rotation-profile dependence. τ is the arrival time of the first peak, while η controls the peak width. This minimal $\{\tau, \beta, \alpha\}$ model achieves high fitting factors against simulation catalogs and provides a direct handle on β while marginalising residual EOS dependence with α [39].

While models like the chirplet and Gaussian triplet offer a phenomenological handle on signal morphology, other techniques such as principal component analysis on waveform catalogues and generative machine learning methods (e.g., GANs [40]) are being developed to overcome the computational cost and stochastic limitations of CCSN modelling. These nonparametric or semi-parametric tools provide alternative ways to generate realistic waveform priors for burst pipelines.

Simplified models

While physically motivated models such as those from compact binary coalescence, cosmic strings, memory effects, and core-collapse supernovae provide deep insights into gravitational-wave emission mechanisms, it is often useful to employ simplified, mathematical waveform models. These allow for tractable approximations of burst-like signals and serve as important tools in the calibration, testing, and development of detection pipelines. Among the most widely used of these are the damped sinusoid and the sine-Gaussian waveforms. The models we show here are in the form they are implemented in `lalsuite` [41, 42].

Damped sinusoids are time-domain models characterised by exponentially decaying oscillations. They serve as toy models for ringdown-like behaviour,e.g., the (f)-mode oscillations of neutron stars or quasi-normal modes from black hole mergers. An elliptically polarized damped sinusoid waveform can be written as:

$$h_+(t) = \frac{1}{2}(1+\cos^2\iota)\,A\,e^{-(t-t_0)/\tau}\cos\left(2\pi f_0(t-t_0)+\phi\right), \tag{24}$$

$$h_\times(t) = \cos\iota\,A\,e^{-(t-t_0)/\tau}\sin\left(2\pi f_0(t-t_0)+\phi\right), \tag{25}$$

where A is the amplitude, τ is the damping time, f_0 is the central frequency, t_0 is the signal onset time, ϕ the phase, and ι the inclination angle.

This waveform captures the qualitative signature of a burst that quickly fades out, and its mathematical simplicity allows it to be used in signal injection campaigns, parameter recovery tests, and waveform discrimination studies. In addition, variants such as those tapered with Tukey windows are frequently employed to avoid edge effects in time series.

Another fundamental waveform model is the sine-Gaussian, typically formulated in the time domain. It models a short-duration signal with a well-defined central frequency and bandwidth, controlled by the quality factor Q. A typical sine-Gaussian strain's Fourier transform is given by:

$$\tilde{h}_+(f) = \frac{h_{\rm rss}\sqrt{\pi}\tau}{2\sqrt{\alpha(1+e^{-Q^2})}}\left[e^{-\pi^2\tau^2(f+f_0)^2} + e^{-\pi^2\tau^2(f-f_0)^2}\right], \tag{26}$$

$$\tilde{h}_\times(f) = \frac{i\,h_{\rm rss}\sqrt{\pi}\tau}{2\sqrt{\alpha(1-e^{-Q^2})}}\left[e^{-\pi^2\tau^2(f+f_0)^2} - e^{-\pi^2\tau^2(f-f_0)^2}\right], \tag{27}$$

where $h_{\rm rss}$ is the root-sum-square amplitude, f_0 is the central frequency, and $\tau = Q/(\sqrt{2}\pi f_0)$ characterises the burst duration.

Sine-Gaussians serve as generic models for narrowband bursts and approximate the response of detectors to a wide class of short-duration events. They are particularly useful in data analysis as they are localized in both time and frequency, and can be parametrically varied across a large portion of the detector's bandwidth.

The utility of these models lies not in their physical realism but in their mathematical generality and ability to mimic real signals. Many physical waveforms, such as those from CCSNe or cosmic-string cusps, can be approximated by sine-Gaussians or damped sinusoids. They provide a basis set for reconstructing unknown signals and serve as a foundation for nonparametric search algorithms. Furthermore, these simplified models allow the probing of the sensitivity of pipelines and evaluate the detectability and distinguishability of more complex waveforms. Because they are flexible and computationally inexpensive, they are commonly used in injection campaigns to benchmark search performance and validate statistical inference techniques [43, 44].

1.2 Gravitational wave data analysis

Gravitational wave data analysis bridges the theoretical predictions of general relativity with the practical challenges of extracting weak astrophysical signals from noisy interferometric data. The process begins with the detection of potential signals amidst instrumental noise and proceeds through successive stages of signal processing, statistical inference, and model selection. This section introduces the key techniques that enable this process: matched filtering for signal discovery, Bayesian inference for model comparison and parameter estimation, and modern computational tools like `bilby` [45, 46] that implement these techniques efficiently. Please refer to [6] for a more complete description of GW data analysis theory.

1.2.1 Signal processing

Gravitational wave detection is fundamentally a problem of signal processing, where we extract weak signals buried in detector noise. Detectors such as LIGO and Virgo are designed to measure extremely weak strains in spacetime, typically on the order of $h \sim 10^{-22}$ or smaller. These measurements are contaminated by instrumental and environmental noise sources, which are best characterised statistically. The data stream from a single detector is modelled as a time series, $s(t)$, a superposition of signal, $h(t)$, and noise, $n(t)$ (Eq. 11). In gravitational wave data analysis, we assume $n(t)$ to be a stationary, Gaussian random process with zero mean. Its statistical properties are fully encoded in its one-sided power spectral density (PSD) $S_n(f)$, defined via the Fourier transform $\tilde{n}(f)$ as:

$$\langle \tilde{n}(f)\tilde{n}^*(f')\rangle = \frac{1}{2}S_n(f)\delta(f-f'), \tag{28}$$

which implies that different frequency components of the noise are uncorrelated and that the variance of the Fourier amplitude at frequency f is proportional to $\sqrt{S_n(f)}$.

In parallel to understanding the statistical structure of the noise, we must also account for the fact that gravitational waves have a tensorial nature, characterised by two polarizations, $h_+(t)$ and $h_\times(t)$. The response of a detector depends on the source's sky location and polarization, encoded in the antenna pattern functions $F_+(\theta,\phi,\psi)$ and $F_\times(\theta,\phi,\psi)$, where (θ,ϕ) are the source's celestial coordinates and ψ is the polarization angle. These functions describe how sensitive a particular detector is to each polarization mode from a given direction. The measured strain

in a detector is therefore:

$$h(t) = F_+(\theta, \phi, \psi)\, h_+(t) + F_\times(\theta, \phi, \psi)\, h_\times(t). \tag{29}$$

In a network of detectors, each interferometer has its own set of antenna pattern functions, and the combined response enables reconstruction of the signal's sky position and polarization content.

To detect and characterise gravitational waves within noisy data, we require a framework to quantify how closely a template waveform $h(t;\boldsymbol{\theta})$, depending on physical parameters $\boldsymbol{\theta}$, matches the observed data $s(t)$. A fundamental concept is the inner product between two time series $a(t)$ and $b(t)$, defined in the frequency domain with PSD weighting as:

$$\langle a|b\rangle = 4\,\mathrm{Re}\int_0^\infty \frac{\tilde{a}(f)\tilde{b}^*(f)}{S_n(f)}df. \tag{30}$$

This inner product naturally induces a norm, $\|a\| = \sqrt{\langle a|a\rangle}$, and allows for the definition of the signal-to-noise ratio (SNR) as:

$$\rho = \frac{\langle s|h\rangle}{\sqrt{\langle h|h\rangle}}. \tag{31}$$

The matched filter SNR is a normalised projection of the data onto the template and is maximized when the template closely matches the true signal. This means that the matched filter SNR has a time dependency (akin to sliding a template unto the data until the best match is found), which is maximed over. In the ideal case where the signal is known exactly and the data contains only signal, the optimal SNR reduces to (and phase dependency is one):

$$\rho^2_{\mathrm{opt}} = \langle h|h\rangle = 4\int_0^\infty \frac{|\tilde{h}(f)|^2}{S_n(f)}df. \tag{32}$$

This integral shows that the SNR is weighted according to the detector's noise properties; frequencies where $S_n(f)$ is small contribute more significantly to the detection statistic. Parseval's theorem ensures that energy is preserved under the Fourier transform:

$$\int_{-\infty}^{\infty} |s(t)|^2 dt = \int_{-\infty}^{\infty} |\tilde{s}(f)|^2 df, \tag{33}$$

justifying the use of frequency domain representations in signal analysis. Extending this to a network of detectors, the total network SNR is computed by summing in quadrature the individual detector SNRs:

$$\rho^2_{\mathrm{net}} = \sum_i \rho_i^2, \tag{34}$$

assuming uncorrelated noise between detectors, which is a generalisation of the inner product with phase term accounted for ($\rho_i \neq \rho_{\rm opt}$). This multi-detector combination significantly improves detection confidence and parameter estimation by leveraging the geographical and orientation diversity of the network.

Matched filtering and coherent detection of modelled signals

This formalism sets the stage for constructing an optimal detection strategy: matched filtering. This technique is derived from the Neyman-Pearson lemma, which states that the most powerful test to distinguish between two hypotheses, signal plus noise $\mathcal{H}_1$ and noise only $\mathcal{H}_0$, is one that maximises the likelihood ratio. Assuming Gaussian stationary noise, the likelihood of observing a particular data stream $s(t)$, given a template signal $h(t;\boldsymbol{\theta})$, is:

$$\mathcal{L}(s|h(\boldsymbol{\theta})) \propto \exp\left(-\frac{1}{2}\langle s-h(\boldsymbol{\theta})|s-h(\boldsymbol{\theta})\rangle\right). \tag{35}$$

The likelihood ratio is then maximised when the template waveform best matches the signal in the data. Expanding the exponent leads to:

$$\log\mathcal{L}(s|h) = -\frac{1}{2}\langle s|s\rangle + \langle s|h\rangle - \frac{1}{2}\langle h|h\rangle. \tag{36}$$

This expression clearly shows that the likelihood is maximised when the inner product $\langle s|h\rangle$ is large, and hence motivates the matched filtering statistic (with a phase term included), Eq. 31, sometimes called the matched filter SNR. When a template exactly matches the signal, ρ is maximised over the phase, and its square corresponds to the optimal signal-to-noise ratio. This derivation confirms that matched filtering is the optimal detection technique in Gaussian noise, assuming the waveform is fully known.

However, the effectiveness of matched filtering is contingent on having accurate waveform models. For CBCs, templates are available that span the full inspiral-merger-ringdown evolution, allowing efficient searches via template banks. But for many other sources, such as cosmic strings, core-collapse supernovae, or unknown astrophysical transients, no such reliable templates exist. In these cases, alternative search strategies are necessary.

To identify generic signals in noisy data, we have "burst pipelines" based on "coherent" and "coincidence" analyses. These pipelines form the front line of detection for poorly modelled signals. One such example is the `Coherent WaveBurst`

(cWB) [47–49] algorithm, a widely used search method that does not rely on templates. Instead, cWB performs a wavelet transform of the detector data into a time-frequency representation using the Wilson-Daubechies-Meyer basis. In this domain, excess power clusters are searched for that appear simultaneously (in time and frequency) across multiple detectors. A signal is considered a viable candidate if it has significant energy in the coherent channels and negligible energy in the null channel. This coherent approach allows one to exploit the consistency conditions across the detector network without making strong assumptions about the waveform morphology.

Once such a candidate event is identified, it transitions into a new phase of analysis: interpretation. In this stage, matched filtering becomes useful again. But now, it is no longer used to detect whether a signal is present, but rather to compare different physical models for the origin of the signal and to infer its parameters. This approach naturally aligns with Bayesian inference, where the likelihood of the data is computed under various waveform hypotheses and the evidence for each is evaluated, as we shall see in the next section.

This two-tiered structure, coherent signal discovery followed by Bayesian model discrimination, encapsulates the traditional logic of burst searches in gravitational wave astronomy. It reflects the epistemological constraints of the field: while some signals are well-modelled and amenable to traditional template matching, others are only constrained statistically through minimal assumptions on their time-frequency behaviour.

1.2.2 Bayesian inference

GW data analysis relies on robust statistical methods to infer the presence of a signal, estimate its physical parameters, and perform model selection. While matched filtering provides an optimal way to search for signals when a well-defined model exists, the problem becomes significantly harder for poorly modelled sources, such as burst signals. A Bayesian approach is particularly well-suited for analysing GW data, as it allows us to incorporate prior knowledge, handle uncertainties systematically, and rigorously compare competing models.

Bayes' theorem

Bayesian inference is based on Bayes' theorem, which updates our belief about model parameters $\boldsymbol{\theta}$ given the observed data d [50, 51]:

$$\mathcal{P}(\boldsymbol{\theta}|d, M) = \frac{\mathcal{L}(d|\boldsymbol{\theta}, M)\pi(\boldsymbol{\theta}|M)}{\mathcal{Z}(d|M)}. \tag{37}$$

Here $\mathcal{P}(\boldsymbol{\theta}|d, M)$ is the posterior probability, describing the updated probability distribution of the parameters given the data; $\mathcal{L}(d|\boldsymbol{\theta}, M)$ is the likelihood, which quantifies how well the model parameters explain the observed data; $\pi(\boldsymbol{\theta}|M)$ is the prior, encoding our a priori knowledge or assumptions about the parameters before observing the data; $\mathcal{Z}(d|M)$ is the evidence, also called the marginal likelihood, which normalises the posterior and is crucial for model selection.

While Bayesian inference is often used for parameter estimation, it also provides a natural framework for model selection. Given two competing models M_1 and M_2, which a priori have probabilities, $\pi(M_1)$ and $\pi(M_2)$, Bayes' theorem gives the posterior odds ratio:

$$O_{12} = \frac{p(M_1|d)}{p(M_2|d)} = \frac{\mathcal{Z}_1(d|M_1)}{\mathcal{Z}_2(d|M_2)}\frac{\pi(M_1)}{\pi(M_2)}, \tag{38}$$

where the ratio of evidences, known as the "Bayes factor" (BF), quantifies how much the data favours one model over the other:

$$\mathcal{BF}_{12} = \frac{\mathcal{Z}_1(d|M_1)}{\mathcal{Z}_2(d|M_2)}. \tag{39}$$

The exact values under which one considers the evidence for a model as opposed to some other as "anecdotal", "moderate", "strong" and "decisive" is a matter of convention. As is common within GW literature, we adopt a $\log_{10}\mathcal{BF} \geq 8$ as "strong" evidence for a model versus another [52].

Bayesian evidence estimation: Nested sampling

The challenge is then to compute the Bayesian evidence:

$$\mathcal{Z}(d|M) = \int \mathcal{L}(d|\boldsymbol{\theta}, M)\pi(\boldsymbol{\theta}|M)d\boldsymbol{\theta}. \tag{40}$$

This integral is high-dimensional and often intractable analytically. MCMC methods are widely used to approximate posterior distributions, drawing samples from $\mathcal{P}(\boldsymbol{\theta}|d, M) \propto \mathcal{L}(d|\boldsymbol{\theta}, M)\pi(\boldsymbol{\theta}|M)$. MCMC functions via iterative Metropolis-Hastings sampling or Hamiltonian Monte Carlo [53–55], but crucially, it does not return

the global evidence $\mathcal{Z}$. Instead, we employ nested sampling, an advanced Monte Carlo technique designed specifically to estimate $\mathcal{Z}$. Nested sampling, introduced by Skilling [56], is a Monte Carlo algorithm explicitly designed to estimate the Bayesian evidence while also providing posterior samples. We utilise the `dynesty` [57] package and its `bilby` implementation.

The key idea behind nested sampling is to convert the high-dimensional evidence integral into a one-dimensional problem by transforming the prior volume:

$$X(\lambda) = \int_{\mathcal{L}(\boldsymbol{\theta})>\lambda} \pi(\boldsymbol{\theta}|M) d\boldsymbol{\theta}. \tag{41}$$

This defines X as the fraction of the prior volume where the likelihood exceeds λ. The evidence integral then rewrites as $\mathcal{Z} = \int_0^1 \mathcal{L}(X) dX$.

The algorithm of nested sampling proceeds as follows:

1. Initialise a set of N live points uniformly sampled from the prior $\pi(\boldsymbol{\theta}|M)$.

2. Identify the lowest likelihood point $\boldsymbol{\theta}_{\rm min}$ in the set and record its likelihood $\mathcal{L}_{\rm min}$.

3. Remove $\boldsymbol{\theta}_{\rm min}$ and replace it with a new sample drawn from the prior but subject to $\mathcal{L} > \mathcal{L}_{\rm min}$.

4. Update the prior volume as it shrinks exponentially:

$$X_k \approx e^{-k/N}. \tag{42}$$

5. Estimate the contribution to the evidence from the removed point:

$$\mathcal{Z} \approx \sum_k \mathcal{L}_k (X_{k-1} - X_k). \tag{43}$$

6. Repeat steps 2-5 until termination criteria (e.g., negligible remaining prior volume) is met, usually $d \log \mathcal{Z} \leq 0.1$.

Nested sampling thus offers several advantages as it directly estimates the Bayesian evidence, naturally handles multimodal likelihood surfaces and it generates posterior samples simultaneously.

1.2.3 Parameter estimation with `bilby`

To implement Bayesian inference in gravitational wave data analysis, we utilise the `bilby` library, a flexible and modular framework designed to facilitate parameter estimation, model selection, and signal injection workflows. Since it is a central cog of BASILIC, we will now outline the typical components of a `bilby` analysis and explain how the toolkit can be used in the context of waveform model comparison for gravitational wave bursts. The `bilby` architecture is centred around several core objects: the *InterferometerList*, the *WaveformGenerator*, the *PriorDict*, and the *GravitationalWaveTransient* likelihood. Together, these components allow one to construct a full inference schema tailored to both modelled and unmodelled gravitational wave signals (Section 2).

Signal injection and noise modelling

For signal injections, the user defines the physical source parameters, such as masses, spins, sky location, polarization angle, and luminosity distance, and converts these into frequency-domain strain data using a specified waveform approximant. This is handled via the *WaveformGenerator* class, which is instantiated with either user-defined source models, or pre-built ones. For CBC signals, this corresponds to *lal_binary_black_hole* with waveform families like IMRPhenomXP. For unmodelled or phenomenological signals (e.g., power-law bursts), users can define their own custom frequency-domain strain functions.

Detector noise is incorporated through an *InterferometerList*, which stores data for each detector in the network (e.g., H1, L1, V1). These detectors are initialised with analytic or tabulated power spectral densities and are capable of simulating noise or reading in real data. For injection studies, signals are added to synthetic noise realisations using *inject_signal*, enabling a realistic test-bed for inference.

Priors and custom models

The prior distribution over the parameters is specified via a *PriorDict*, which supports standard forms (Uniform, LogUniform, DeltaFunction, etc.) and astrophysically motivated distributions (e.g., uniform in comoving volume or cosine of inclination). When comparing models with different physical assumptions, such as a binary black hole versus a generic burst signal, each model is assigned its own

prior dictionary reflecting its parameter space.

Bayesian inference and nested sampling

The likelihood function used for transient signals is *GravitationalWaveTransient*, which computes the standard frequency-domain Gaussian likelihood, marginalised over antenna pattern response and including both polarizations. The priors and likelihood are passed to the *run_sampler* routine, where nested sampling is employed for evidence estimation and posterior generation. In our case, we employed *dynesty* with advanced tuning parameters, e.g., live points, acceptance thresholds, and parallelization settings, to improve convergence and handle multi-modal posteriors.

For each model, the sampler returns posterior distributions over parameters and an estimate of the $\log \mathcal{Z}$. The difference in evidences between models defines the BF. When performing injections, the maximum-likelihood (ML) parameters are extracted from the posterior samples to allow reinjection and visualisation of recovered waveforms.

The `bilby` workflow enables a comprehensive Bayesian approach to parameter estimation and model selection in gravitational wave data analysis. By supporting both predefined waveform families and user-specified models, it facilitates a direct comparison between competing hypotheses, through the evidence computed via nested sampling. This library allows finely tuned analyses of burst GW signals, but it does require extensive user knowledge, of all the topics discussed so far, as well as coding, high performance computing and others. To unlock this utility—not only the recovery of source parameters, but also robust statistical discrimination among models, a critical step in the interpretation of gravitational wave bursts with uncertain origin—to a wider scientific audience, we have built Basilic, a no-coding, user-friendly GW burst signal analysis pipeline, the topic of the next chapter.

1.3 Conclusion

In this chapter, we have covered most of the relevant theoretical aspects of gravitational wave in general and burst signal sources in particular, as well as the data analysis methods currently available in their study. The inclusion of both complex physical waveforms and simplified phenomenological models in gravitational-wave

analyses highlights the fundamental challenge of burst signal detection and classification. Unlike well-modelled signals from compact binary inspirals, burst signals often:

- Lack precise theoretical predictions,
- Exhibit significant morphological diversity,
- Arise from disparate physical processes,
- Show partial overlaps in spectro-temporal features.

Thus, the detection and interpretation of burst signals represent one of the most intellectually rich and technically challenging aspects of gravitational-wave data analysis. The combination of simplified and complex waveform models provides the necessary toolbox to navigate this landscape and extract meaningful physics from short-duration gravitational-wave events.

In the next chapter, we will present the major output of this work, BASILIC. This is a pipeline constructed specifically to streamline the process of data analysis of burst signals, and more importantly, make it accessible for any GW scientist to tackle the analysis, without delving into actual coding. Furthermore, by automating the entire process, from PSD and data segment selection to output, BASILIC unlocks a large class of previously sparsely-tackled problems in GW science, injection studies. This is achieved through an integrated module of burst signal models that allows rapid hypothesis testing on injected data.

2 Basilic

The LVK workflow for parameter estimation is usually organised around one event at a time, a natural consequence of the humble beginnings of LVK sensitivity curves, with a consequently modest event detection rate in the first two observing runs. As this improves (two CBCs per week expected during the fourth observing run), we are slowly moving towards a future of multiple alerts by online pipelines daily, and countless other events that are not identified on the fly, but could prove interesting in offline studies. After a candidate is identified by an online search, detailed parameter estimation follows, often with bespoke scripts, partly manual bookkeeping and interactive post-processing. This pattern works for compact binary coalescences, where mature tools exist and the hypothesis space is narrow, but it scales poorly to short-duration burst sources, since they are heterogeneous. Different physical scenarios can lead to similar time-frequency morphologies, as mentioned in the last chapter, especially at modest signal-to-noise ratio. It is entirely plausible that real signals remain under-classified by low-latency searches, or that interesting sub-dominant hypotheses go untested without a targeted offline analysis. Furthermore, access to these models is also a limiting factor, with `bilby` requiring manual implementation, testing and validation.

The barrier here is not just computing time and model implementation. It is also the coding needed to wire said models, priors, samplers and plotting into a robust workflow. `Bilby` already provides the statistical machinery, but it expects the user to write analysis code and to integrate domain-specific models. This is why `bilby-pipe` was built, but it is centred around the longer duration generic CBC events. Similarly, this is precisely what Basilic addresses for bursts. The goal is a pipeline that requires no user coding: the user specifies a study in a short configuration file, submits jobs to the cluster, and receives a clean summary with the essential plots and tables. With friction removed, it becomes practical to test alternative burst hypotheses for any event of interest and to run large injection studies that probe sensitivity, parameter recovery and model confusion, all with minimal user intervention.

In this chapter, we will provide a description of the main body of work of this thesis, the creation of Basilic, its philosophy, currently available features and implemented models, the user experience, as well as our trajectory for the future.

All the results presented in the next chapter, have been obtained with the use of BASILIC on the LVK HTC cluster at Caltech.

2.1 Design philosophy

BASILIC, an acronym for BAyesian tranSIent modeLIng and Classification is built on the following guiding principles:

- **User-facing simplicity.** A single `YAML` file describes the study: which data to analyse, which PSD to use, which models to test, how to sample, how many injections if requested. The rest is automated.

- **Leverage, not replace, `bilby`.** We call `bilby` for likelihoods, waveforms and sampling. BASILIC adds orchestration, a growing model registry, I/O and post-processing.

- **HTCondor-centred execution.** The pipeline ships ready to submit on the CIT HTCondor pool. Local runs are also supported. Injection campaigns map to a DAG with batched jobs.

- **Reproducible outputs.** We keep the native `bilby` result directory intact and write a second, curated output with plots, CSV summaries and an HTML report. Identical configuration files produce statistically identical output.

- **Extensibility.** Models are registered in a central provider with internal names and display labels. Adding a new model is a small, local change.

Table 1 lists the waveform families currently provided by the `GWModelProvider` module. For each entry we note the internal model name used in the configuration, the evaluation domain (time or frequency), the intrinsic parameters expected by the source model, and a brief description. Source extrinsic parameters, right ascension, declination, phase and geocentric time $\{\mathrm{ra}, \mathrm{dec}, \psi, \mathrm{t\_geo}\}$ are included in the prior and passed through to the likelihood.[1] Each parameter definition has been given previously in Sec. 1.1.3 along with the equations for the waveform polarizations for all non-CBC models shown and references.

[1] Compact-binary approximants follow the standard `bilby`/`LAL` parametrisation; we do not repeat it here, it is a 15-dimensional parameter space for BBH CBC, including masses, spins etc.

Table 1: Waveform models currently implemented in BASILIC. Model names in the first column are exactly the strings used in the configuration file. Intrinsic parameters are the internal keywords; display labels used in plots are set in the model registry.

Model name	**Domain**	**Parameters**	**Notes**
`asymmetric_gaussian_chirplet`	time	A, f, $\dot{f}$, c, Q, t_0, a	Eqs. 21, 22.
`ccsn_gaussian_triplet`	time	α, β, τ, d	Eq. 23.
`generic_power_law`	freq.	A, $f_{\rm low}$, $f_{\rm high}$, α	Eq. 19.
`cosmic_string_cusp`	freq.	A, $f_{\rm low}$, $f_{\rm high}$	Eq. 17.
`cosmic_string_kink`	freq.	A, $f_{\rm low}$, $f_{\rm high}$	Eq. 18.
`cosmic_string_kink_kink`	freq.	A, $f_{\rm low}$, $f_{\rm high}$	Ref. [28]
`memory_effect`	freq.	A, $f_{\rm low}$, $f_{\rm high}$, τ	Eq. 20.
`sine_gaussian`	freq.	h_{rss}, τ, f_0, α	Eq. 26.
`damped_sine`	freq.	A, τ, f_0, t_0, ϕ, ι	Eq. 24.
`IMRPhenomPv2`, `IMRPhenomD`, `IMRPhenomXP`, `IMRPhenomXPHM`, `NRSur7dq4`	freq.	Standard CBC	Ref. [41]

Following the first design principle, BASILIC provides a lot of output for very little input. The entirety of the mandatory work of the user, is a single configuration file describing the analysis to be done. Optionally, the user can later choose to modify the default priors for each model. As output, BASILIC writes two complementary trees: `result_bilby/`, one subdirectory per model with the native `bilby JSON` result, sampler logs and diagnostics; `result_final/`, curated outputs. We save per-hypothesis CSV files with SNRs and sampler diagnostic columns, a master CSV with one line per model (log evidence, BF, maximum-likelihood and maximum-posterior parameter sets), plots (PSD, SNR histograms, corner plots and more) and an HTML summary. A data-centric flow chart can be found in Fig. 4 in summary of the aforementioned.

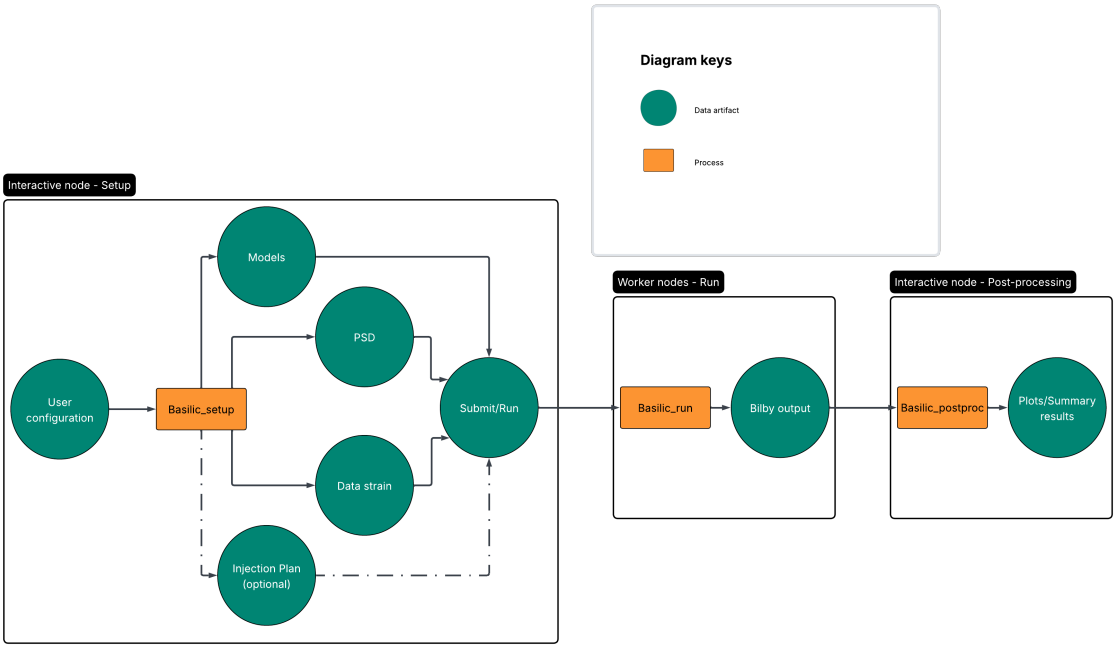

Figure 4: BASILIC data-flow. A `YAML` configuration drives the pipeline. On the interactive node, BASILIC_SETUP validates inputs, specifies data and PSD, models, and emits HTCondor submission files; if enabled, an injection plan (dashed line) augments the DAG. On worker nodes, BASILIC_RUN performs Bayesian inference with pre-built burst models, producing per-sample posteriors (CSV), `bilby JSON`, and SNR CSVs. Finally, on the interactive node, BASILIC_POSTPROC generates plots and an HTML summary plus a machine-readable `summary_results.csv`.

2.2 User experience

From the user's point of view there are three scripts: BASILIC_SETUP, BASILIC_RUN and BASILIC_POSTPROC. The pipeline can handle different data and PSD sources alongside the hypothesis to be tested, as well as injection campaigns, all generated by the first script:

1. **Data.** One of: (i) GWF files already on disk; (ii) a channel and time interval to download with `gwpy`; (iii) Gaussian noise generated from a chosen PSD. For injection studies the same mechanism is used, with the injection applied deterministically on top of each realisation.

2. **PSD.** One of: (i) a `.dat` file on disk; (ii) an estimate computed from a user-provided data segment; (iii) design sensitivity via the built-in `bilby` curves. For Gaussian noise we can use the analysis PSD or a separate PSD dedicated to noise generation.

3. **Hypotheses.** A list of model names as in Table 1, with priors read from `YAML` files. Multiple models can be tested in one go, with consistent extrinsic priors.

After deployment of the pipeline on an interactive node of the CIT cluster, the setup script materialises data and PSDs on disk and writes HTCondor submission files. For single-event analyses we submit a single job per model. For injection campaigns we write a DAG that groups many noise realisations per compute job, with one injection definition applied uniformly. Seeds are handled deterministically so that runs are reproducible and per-detector realisations remain independent.

The run script calls `bilby` and writes results to `result_bilby`. When all jobs are complete, the post-processing script gathers all of this and generates the plots and the HTML summary. The abbreviated HTML embeds a few key plots as `base64` so it can be shared as a single file. The full HTML references the image files on disk and adds heavier plots.

In summary, the entire user experience consists of installing the pipeline, writing configuration file and running the setup script, optionally modifying priors, submitting the pre-built DAG to HTCondor, and running the post-processing script. This simple schema results in BF between models, posterior corner plots and SNR per model, maximum a posteriori waveform comparison between hypothesis in the time domain as well as spectrograms, and more. As a final note, these three scripts have not been merged into one automatic step, because a) it is recommended that the user checks the data and PSD plots before submitting to HTCondor (computational time is precious), and b) some users might not need result post-processing or wish to handle it manually.

2.3 Ongoing and future development

The implementation already covers end-to-end studies for burst hypotheses and compact binaries, with HTCondor integration and injection campaigns. There are three active development threads:

- **Model coverage.** We plan to add more parametrised bursts that are used in practice, and to connect to non-parametric reconstructions for comparison.

- **Trans-dimensional exploration.** Today we test one model per sampler run. A natural extension is reversible-jump sampling over a set of burst models, so

that evidence ratios and parameter inference are performed in a single trans-dimensional chain.

- **Campaign tooling.** Additional summaries for large injection sets, including efficiency curves, confusion matrices between model families and automated quality check of outliers.

- **Going public.** After a more feature rich version is reached, the aim is to make the pipeline available to the broader scientific community.

The source is hosted on the LVK GitLab [58], available to all LVK members. The repository includes example configurations, a parameter reference and usage guides, with defaults chosen to enable a first run without edits. The intent is to keep the user interface stable, even as the model registry and the internals grow.

In the next chapter, we apply BASILIC to two concrete use cases: re-analysis of a recent BBH event under alternative burst hypotheses, and an injection campaign that probes recoverability and model confusion across a range of SNR and morphologies.

3 Use cases and results

This chapter shows how BASILIC is applied in practice. The aim is twofold. First, we demonstrate that a user can take a real event and test more than one physical hypothesis with minimal effort. Second, we show how the same framework scales to studies that require many runs, such as injection campaigns. The emphasis here is on the scientific use-case, rather than on software details. The workflow follows the flowchart, shown in Fig. 4. We begin with a single-event analysis of GW231123 and then move to controlled experiments with simulated signals.

3.1 Individual event analyses - GW231123

On November 23, 2023, at 13:54:30.634 UTC, the two LIGO detectors in Hanford, Washington and Livingston, Louisiana detected a brief gravitational wave event. GW231123 [4] is a short signal with power concentrated in a narrow time-frequency patch. GW231123 was chosen as an object of study to complement the main LVK paper [4]. It assumes a BBH origin of the signal, under this hypothesis, they find the largest yet pair of stellar mass black holes, with source frame masses of 137^{+22}_{-17} $M_\odot$ and 103^{+20}_{-52} $M_\odot$ (90% credible intervals).

GW231123 resembles other heavy-mass CBC candidates in the catalogue where only a handful of cycles enter the detector band. A well known predecessor is GW190521 [59, 60]. That event contained around 5 signal cycles, and was analysed with compact-binary coalescence models as well as with burst signals, with a frequentist method. The purpose was to test whether a short and loud feature could be explained by a cosmic-string (CS) cusp or by an eccentric or high-mass binary black hole.

Both of these events constituted, at the time of detection, the largest stellar BBH system discovered. Yet, the collaboration paper on GW231123, does not consider other origin sources. This remarkable event warrants then further analysis and origin confirmation, within a fully Bayesian setting, complete with parameter estimation and hypothesis ranking. The study presented below was carried out with BASILIC and is written up in a dedicated manuscript, currently under review at *Physical Review Letters* [5]. The work presented here was not part of the main collaboration manuscript due to the analysis being carried out late in the internal review process.

3.1.1 Data and hypotheses

The analysis here uses the same data segment for GW231123, as in Ref. [4], with identified narrow features removed. Two broad hypotheses are tested. The first assumes a BBH compact binary coalescence described by a state-of-the-art precessing surrogate waveform, `NRSur7dq4` [61], which is well-suited to high-mass, potentially highly spinning binaries. The second assumes a linearly polarized burst from a cosmic string. Three emission channels are considered: cusp, kink, and kink-kink collisions. In addition, a generic frequency-domain power-law burst is included to capture a class of short signals that need not follow any of the fixed CS spectral indices. The comparison is performed jointly across the two LIGO detectors in observation at the time of the event.

The Bayesian method is described in Sec. 1.2.2, and the BF for model comparison reported are relative to the no-signal hypothesis. Furthermore, the prior odds for all models are considered equal. The CS models can be found in eqs. (17) to (19). In order to robustly compare BF, we reproduce the analysis carried out [4] both in terms of BBH CBC parameter priors, as well as utilising the same cleaned PSD, with our analysis being run entirely with BASILIC, rather than `bilby_pipe`. As for the CS model priors, we scan a parameter space as large as possible, while remaining computationally feasible. The amplitude is LogUniform from 10^{-23} to 10^{-18} while the high frequency cut-off is allowed to vary uniformly between 25 Hz and 2000 Hz, where the sensitivity drops off. The low frequency cut-off matches the detector's sensitivity decrease at 20 Hz. Lastly, the temporal localization matches [4], while the sky localization prior is maximal. As for the spectral index prior, we will present three different analyses with different priors, since this is the only parameter where prior selection is both arbitrary and can impact the posterior. We consider uniform priors, with an upper limit of 4 for the spectral index (highest index of cosmic strings is 2), while the lower limits of the three priors considered are 0.01, 0.1 and 0.75 (as the spectral index decreases, the signal starts to resemble a straight noise-like line in the time domain).

3.1.2 Results

Assuming a binary black hole origin and employing the `NRSur7dq4` approximant, we obtain posterior distributions that agree with the LVK analysis of GW231123 [4],

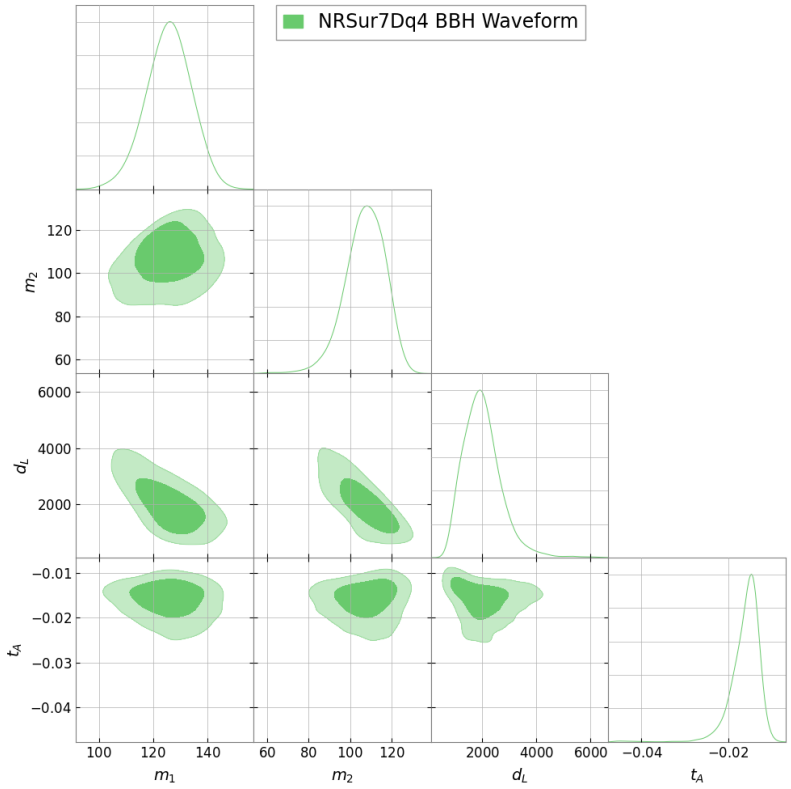

Figure 5: A subset of intrinsic, source-frame parameters (component masses), as well as some extrinsic parameters (luminosity distance and relative-to-trigger arrival time) are shown in green. These are obtained for GW231123 using the `NRSur7dq4` BBH approximant. The coloured contours represent 1-σ and 2-σ credible regions.

shown in Fig. 5. The preferred configuration is a coalescence with component masses $m_1 = 125.98^{+13.5}_{-14.6}\,\mathrm{M}_\odot$ and $m_2 = 107.77^{+14.1}_{-17.3}\,\mathrm{M}_\odot$ at a luminosity distance $d_L = 1932^{+1339}_{-952}$ Mpc. This model yields the largest network SNR and the strongest support relative to noise, as summarised in Table 2 together with the BF for all tested hypotheses. These estimates are consistent with those reported by the LVK using `NRSur7dq4` alone, namely $m_1 = 129^{+15}_{-14}\,\mathrm{M}_\odot$, $m_2 = 111^{+14}_{-17}\,\mathrm{M}_\odot$, and $d_L = 1500^{+1500}_{-800}$ Mpc [62].

For the CS scenarios, the cusp, kink, and kink-kink collisions favour high-frequency cut-offs of $f_{\mathrm{high,\,cusp}} = 67.3^{+1.6}_{-2.6}\,\mathrm{Hz}$, $f_{\mathrm{high,\,kink}} = 67.6^{+1.6}_{-2.8}\,\mathrm{Hz}$, and $f_{\mathrm{high,\,kink-kink}} = 67.9^{+3.2}_{-3.2}\,\mathrm{Hz}$, respectively. These posteriors, shown in the left panel of Fig. 6, are compatible with the merger frequency of a binary black hole with total mass of order $200\,M_\odot$. The same behaviour appears in the power law analysis (right panel of Fig. 6). Among the non-binary black hole hypotheses, the power law yields the largest BF relative to noise. This outcome follows from the additional freedom in

Table 2: The $\log_{10}\mathcal{BF}$ between the waveform model considered and the noise (no-signal) hypothesis, as well as network matched filtering signal to noise ratios are reported below. The values for the power law correspond to the Uniform $(0.75, 4)$ spectral index prior analysis.

Waveform Model	$\log\mathcal{BF}$	SNR
Binary Black Hole	82.13	$21.01^{+0.2}_{-0.3}$
Cusp	38.02	$14.65^{+0.1}_{-0.3}$
Kink	30.91	$13.59^{+0.1}_{-0.3}$
Kink-kink	24.40	$12.41^{+0.2}_{-0.4}$
Power law	45.23	$16.16^{+0.1}_{-0.3}$

the spectral index, which the data favour at lower values, offsetting the broader prior volume compared with the fixed-index cusp, kink, and kink-kink models. The right panel of Fig. 6 shows that changing the prior on the spectral index produces posteriors that are largely prior driven, with a clear preference for the lowest allowed values; this tendency also accounts for the relative preference of the cusp over the kink case.

The discrepancy between the CS templates and the data is apparent in the time domain, shown in Fig. 7. In that figure we show the maximum a posteriori for the cusp model against the detector strain, together with the 90% credible region for the `NRSur7dq4` binary black hole waveform. Because the CS spectra enforce a power-law shape, the corresponding time-domain signals span only about three cycles. These few cycles match roughly half of the observed structure in the strain and lead to a bimodal posterior for the arrival time. The same behaviour appears for the cusp and kink hypotheses (more weakly for the kink), and it is most evident for the generic power-law model with a free spectral index. In contrast, the binary black hole model shows no such bimodality in arrival time, as can be seen in Fig. 5.

We emphasize that both the CS and power-law models achieve SNRs above the typical threshold of ~ 8 and give strong evidence for a signal over noise. However, when compared directly with the best-fitting binary black hole waveform, they are strongly disfavoured. This underlines the importance of performing model comparisons across several plausible hypotheses when classifying short-duration events.

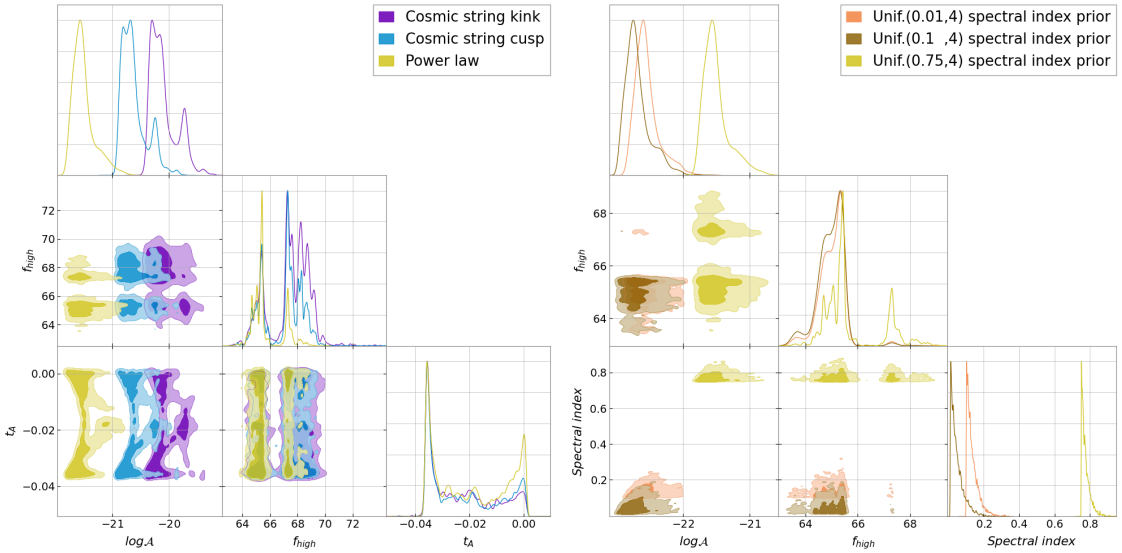

Figure 6: Corner plots of posterior samples, for three models considered. Left, CS kink (purple), CS cusp (blue) and arbitrary power law (yellow), the log-amplitude, high frequency cut-off and relative-to-trigger arrival time are shown, with 1-σ and 2-σ contours. The kink-kink model is also considered and consistent with the above posterior, but not shown here for brevity. Right, the arbitrary power law's spectral index is considered with three different uniform priors (starting at 0.01 (orange), 0.1 (brown) and 0.75 (yellow) respectively), the latter coinciding with the power law on the left. The extra freedom allowed with the power-law model shows clear prior-dominated posteriors.

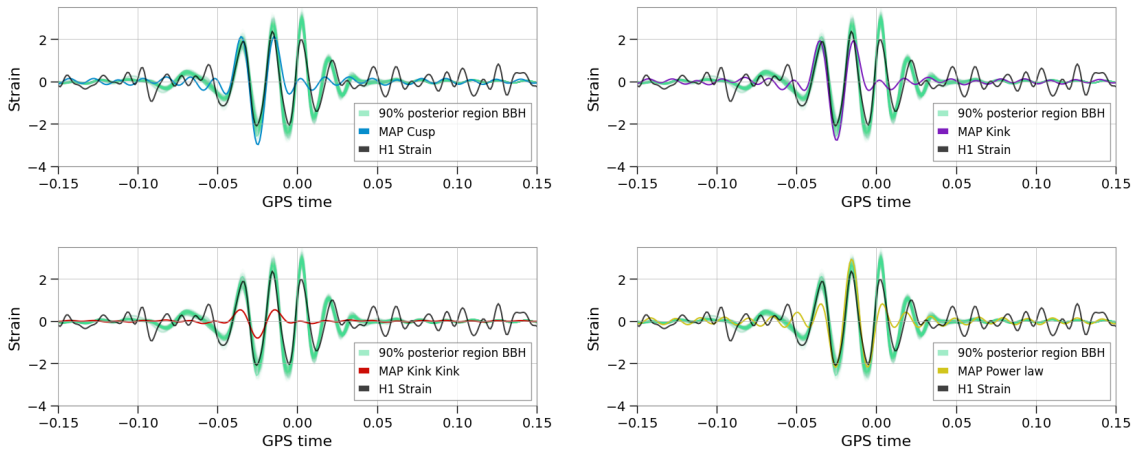

Figure 7: Best-fitting CS cusp waveform (blue) and the LIGO Hanford detector strain (black) are shown in the upper left plot. Upper right is the same for kink waveform (purple), lower left is the kink kink waveform (red) and lower right is the power law with 0.75 lower limit on the spectral index prior (yellow). All data in these plots is whitened with the BayesWave PSD and band-passed between 20-128 Hz. The 90% confidence region of the NRSur7dq4 binary black hole waveform is superimposed (green).

3.1.3 Conclusions

We have shown a Bayesian model-comparison study to assess the origin of GW231123 [4]. Alongside a compact-binary coalescence hypothesis modelled with `NRSur7dq4`, we tested the three principal CS emission channels (cusp, kink, and kink-kink) and, motivated by the power-law character of their spectra, a generic power-law burst with a free spectral index. The data favour the binary-black-hole interpretation over each CS alternative and over the generic power law. This outcome echoes the analysis of GW190521 [59], where a Bayesian comparison likewise preferred a binary-black-hole scenario over CS cusps [60].

Both GW190521 and GW231123 exhibit only about five cycles in band. Future detections may feature even shorter signals. Systematic comparisons across plausible short-duration hypotheses will therefore remain essential for robust source classification. With Basilic, such studies become routine. The configuration used for GW231123 is provided in Appendix A, together with brief notes on the fields that control the analysis window, PSD estimation, and model list. This same workflow extends naturally to large injection campaigns and broader surveys of competing burst morphologies.

3.2 Injection studies

We now present the outcomes of a controlled injection campaign that we ran with the proto-version of BASILIC built on top of `bilby`. The goal is to test model recovery and model selection for short gravitational-wave transients in a setting where the ground truth is known. We generate synthetic data by adding a specified signal to stationary, Gaussian, coloured noise constructed from the design sensitivity curves of LIGO Hanford, LIGO Livingston, and Virgo. Each distinct noise draw constitutes one *injection* for a fixed signal and parameter choice, and a batch of injections at common settings is used to build statistics. Because BF depend on the accuracy of the numerical evidence estimate, single realisations can produce deceptively large BF even when the injected parameters are not well recovered. This can happen when a flexible model happens to match idiosyncrasies of one noise draw. Averaging over many independent realisations mitigates this effect. For any one realisation, a combined view of model comparison (through BF) and parameter-estimation diagnostics (marginals and trace plots) gives the clearest picture of performance.

All examples below use 4 s of data sampled at 4096 Hz. Depending on the study we vary either the injected network SNR or the signal parameters, always over multiple noise realisations. The sky position and extrinsic parameters are fixed across runs to reduce degeneracies: right ascension $\alpha = 1\,\mathrm{rad}$, declination $\delta = 0.85\,\mathrm{rad}$, geocentric arrival time $t_{\mathrm{geo}} = 0\,\mathrm{s}$, and polarization angle $\psi = 1.5\,\mathrm{rad}$. We first show direct recovery of injected waveforms for CS and memory models, then a two-model discrimination test with a CS injection and recovery with both CS and memory templates, and finally a confusion study between high-mass binary black holes and cosmic strings. All runs were orchestrated and post-processed with BASILIC.

3.2.1 Cosmic strings, memory, and parameter recovery

We inject a CS *kink* signal, Eq. 18, into Gaussian noise coloured by the design PSDs of H1, L1, and V1, and we recover with the generic power-law model, Eq. 19 using the priors in Table 3. To focus the inference on the core morphology, we fix the extrinsics as listed above and sample only the amplitude, spectral index, and arrival time. In a parallel set of runs we inject the memory-effect model, Eq. 20, and recover with the same memory template, sampling amplitude, rise time, and arrival

	Parameter [Unit]	Prior Type
Shared Priors	α [rad]	DeltaFunction[mean = 1]
	δ [rad]	DeltaFunction [mean = 0.85]
	$t_{\rm geo}$ [s]	DeltaFunction [mean = 0]
	ψ [rad]	DeltaFunction [mean = 1.5]
Generic Power Law	A_p [N/A]	LogUniform [min = 10^{-22}, max = 10^{-19}]
	α [N/A]	Uniform [min = 0.1, max = 2]
	t_c [s]	Uniform [min = 0, max = 4]
	$f_{\rm low}$ [Hz]	DeltaFunction [mean = 20]
	$f_{\rm high}$ [Hz]	DeltaFunction [mean = 2000]
Memory Effect	A_m [N/A]	LogUniform [min = 10^{-23}, max = 10^{-18}]
	τ [s]	LogUniform [min = 10^{-4}, max = 10^{-1}]
	$t_{\rm arr}$ [s]	Uniform [min = 0, max = 4]
	$f_{\rm low}$ [Hz]	DeltaFunction [mean = 20]
	$f_{\rm high}$ [Hz]	DeltaFunction [mean = 2000]

Table 3: Priors for the generic power-law and memory-effect waveforms used in the injection-recovery tests. A hard low-frequency step at 20 Hz is applied to both models to regularize the $f \to 0$ behaviour, and a high frequency cut-off for the memory model is included to stabilize the likelihood.

time with the priors in Table 3.

Figure 8 shows representative posteriors. In Fig. 8a the power law injection-recovery at network SNR $\simeq 7$ places the injected values well inside the 95% credible regions, consistent with the rule of thumb that parameter recovery stabilizes when the network SNR is close to the detection threshold, conventionally taken as $\rho \simeq 8$. In contrast, in Fig. 8b the memory injection at network SNR $\simeq 4$ is not well constrained: the injected values lie outside 99% credible regions, which indicates that the sampler cannot lock on to the true parameters at this low SNR. This same behaviour appears for other morphologies (BBH, CCSN) when the signal is buried under noise.

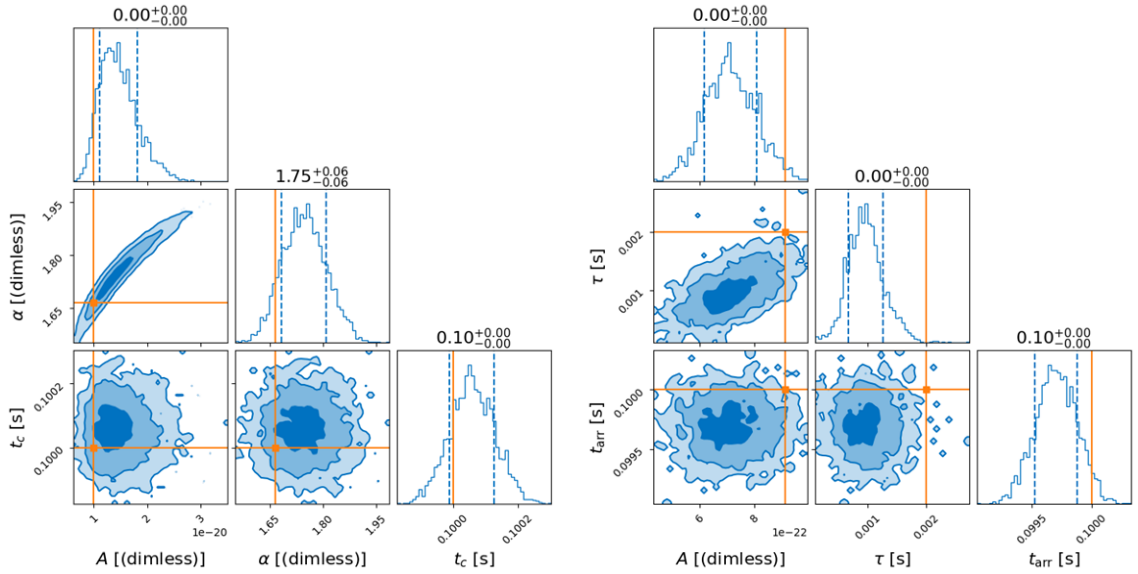

(a) Generic Power Law (b) Memory Effect

Figure 8: Posterior corner plots for two injection-recovery examples. Left: power-law injection at network SNR ~ 7 recovered with the power-law model; the orange lines mark injected values, and 1-σ, 2-σ, 3-σ contours are shown. Right: memory-effect injection at network SNR ~ 4 recovered with the memory model; recovery fails to capture the true parameters at such low SNR.

3.2.2 Bayes factor as a discriminator

We next move to a simple model-selection test. A CS signal is injected at 15 amplitudes that span a range of network SNRs. For each SNR we generate ten independent noise realisations and compute the average BF for two signal hypotheses, cosmic strings and memory effects, each against the noise-only hypothesis. This follows the canonical setup where one asks which of two competing models better accounts for the same data set.

Figure 9 shows the BF as a function of SNR. Both signal hypotheses overtake noise as SNR grows, which is the expected exponential trend from Eq. 36 through the ρ^2 scaling of the log-likelihood in the Gaussian limit. The gap between the two signal curves widens with SNR, which means that the correct model (cosmic strings here) becomes preferred over the alternative (memory) at an increasing rate as the signal strengthens.

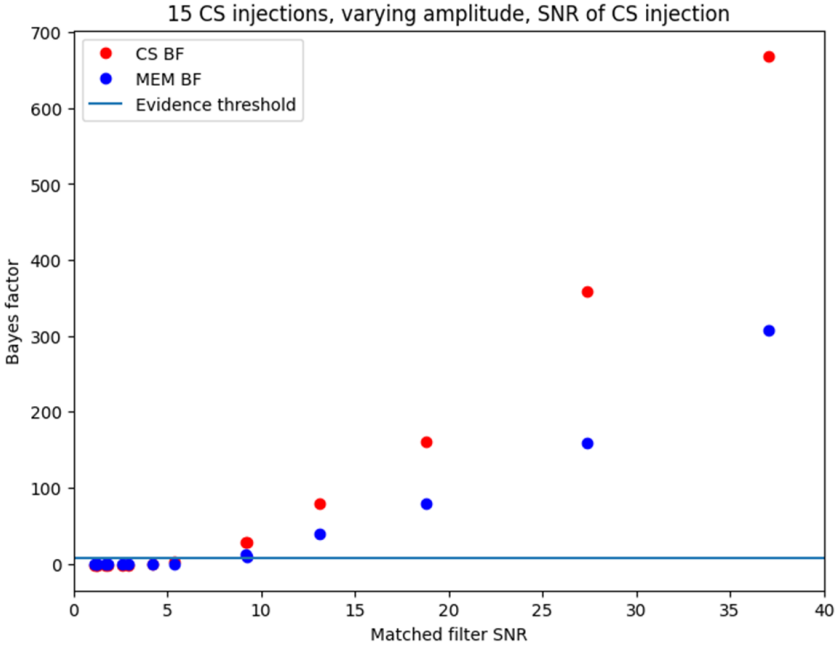

Figure 9: Mean $\log \mathcal{BF}$ against matched filter SNR. Red points: CS signal hypothesis versus noise. Blue points: memory-effect hypothesis versus noise. Each point is the mean over ten noise realisations at fixed SNR. The cyan line indicates $\log \mathcal{BF} = 8$, often used as a "strong" evidence mark.

3.2.3 High-mass BBH and cosmic strings

We now examine a plausible confusion between a well-modelled source and a simple burst morphology. Very massive BBH signals can present only a few cycles in band, which produces a compact time-domain feature that resembles a low-index power law in frequency. To test whether this resemblance translates into model-selection ambiguity we proceed in two steps, keeping the injected network SNR fixed at ~ 6 for comparability. This is exactly the regime that justifies analyses like the one performed for GW231123. Had the SNR been lower, or the source masses higher, a BBH event could be mistaken for a cosmic string or vice-versa. This confusion needs to be quantified, which is our aim here.

First, we inject three equal-mass BBH systems and scale the luminosity distance so that the network SNR remains constant. The systems are chosen to span different durations in band: a lower-mass case with $\sim 60\,M_\odot$ per component, an intermediate case with $\sim 115\,M_\odot$, and a very high-mass case with $\sim 205\,M_\odot$. Each signal is injected into 200 independent noise realisations. For every realisation we perform Bayesian analyses under a BBH hypothesis and a CS hypothesis with the priors in

	Parameter [Unit]	Prior Type
Shared Priors	α [rad]	DeltaFunction[mean = 1]
	δ [rad]	DeltaFunction [mean = 0.85]
	$t_{\rm geo}$ [s]	DeltaFunction [mean = 0]
	ψ [rad]	DeltaFunction [mean = 1.5]
Generic Power Law	A_p [N/A]	LogUniform [min = 10^{-24}, max = 5×10^{-18}]
	α [N/A]	Uniform [min = 0.1, max = 3]
	t_c [s]	Uniform [min = 0, max = 2]
	$f_{\rm low}$ [Hz]	DeltaFunction [mean = 20]
	$f_{\rm high}$ [Hz]	Uniform [min = 80, max = 350]
BBH CBC	m_1, m_2 [$M_\odot$]	Uniform [min = 40, max = 300]
	d_L [Mpc]	Uniform [min = 10000, max = 30000]

Table 4: Priors for the generic power law and BBH CBC recovery in the BBH-vs-CS study. All other BBH intrinsic and extrinsic parameters are fixed (zero-spin system) to reduce computational cost.

Table 4.

Figure 10 shows the outcome. The lower-mass injections (blue) cluster near or below zero on the vertical axis, indicating that a short power-law burst is not a good description; in many realisations the noise hypothesis even outperforms the CS model. This acts as a sanity check because a longer in-band BBH signal should not mimic a brief burst. The BBH $\log \mathcal{BF}$ for these points ranges from about 2 to 10, which matches expectations for a network SNR near 6. For the higher-mass cases the relationship changes: as the BBH BF grows, the CS BF tends to grow as well. In most realisations ($> 90\%$) there is no strong BBH-versus-CS preference, even though BBH wins in pairwise comparisons more often than not. The morphology overlap at high total mass therefore does translate into practical ambiguity in a standard BF plot.

Second, we flip the experiment. We inject two CS *cusp* signals that differ in the high-frequency cut-off, which is the parameter most closely tied to the source geometry. We keep the network SNR at ~ 6 by adjusting the amplitude. For the BBH recovery we sample only the luminosity distance and component masses to keep the computation tractable, and we run 100 noise realisations for each injected

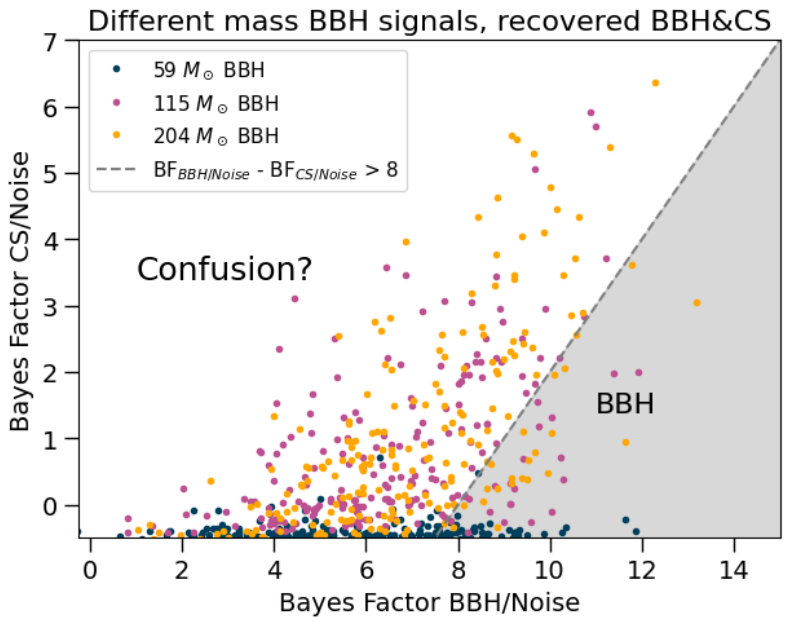

Figure 10: $\log \mathcal{BF}$ versus $\log \mathcal{BF}$ for BBH injections. Horizontal axis: BBH versus noise. Vertical axis: cosmic string versus noise. Each point is one noise realisation. The three colours correspond to equal-mass binaries with component masses indicated in the legend. The shaded region marks cases where BBH is strongly favoured over cosmic strings.

cusp. Figure 11 presents the corresponding BF scatter. Only a modest fraction of realisations ($\sim 15\%$) lead to a clear identification of the data as a CS signal. The same tendency observed in the BBH-injection study appears here as well, although less strongly: when the CS hypothesis looks better than noise, the BBH hypothesis often also improves.

3.2.4 Conclusion

The injection campaign clarifies three points. First, parameter recovery stabilizes once the network SNR approaches the usual detection threshold, while at lower SNR the posteriors become unreliable regardless of morphology. Second, BF against noise scale as expected with SNR and correctly favour the injected model over close competitors. Third, there exist genuine confusion regimes in which very high-mass BBH signals and low-index power-law bursts earn comparable support from the same data, which motivates multi-hypothesis analyses for short transients. All runs in this section were orchestrated with Basilic, which automated the generation of Gaussian, coloured noise realisations, deterministic seeding per interferometer,

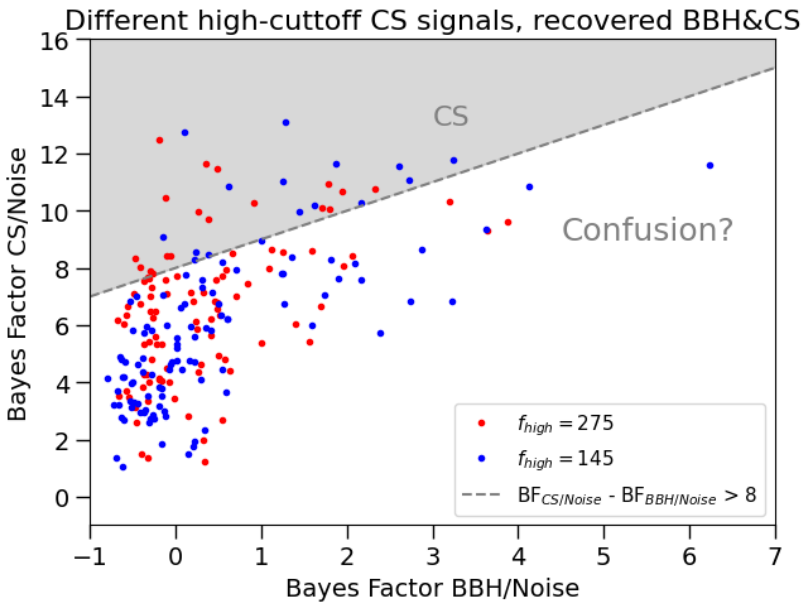

Figure 11: $\log \mathcal{BF}$ versus $\log \mathcal{BF}$ plot for CS cusp injections. Horizontal axis: BBH versus noise. Vertical axis: cosmic string versus noise. Two cusp signals with different high-frequency cut-offs are injected; each point is one noise realisation. The shaded area marks strong preference for cosmic strings.

DAG-based batching on HTCondor, and standardized post-processing. This made the execution of hundreds of realisations routine and let us focus on the scientific interpretation and on planning follow-up studies. Chronologically, the BBH-versus-CS confusion tests were carried out before the GW231123 analysis and served as a practical motivation for developing and using BASILIC; the single-event study then illustrates how the same toolset applies cleanly to public data. The work quantifying the BBH vs CS confusion in real detectors is part of a manuscript currently in preparation.

Conclusion and discussion

This thesis set out to lower the practical barrier to hypothesis-aware inference for short duration gravitational wave transients. The result is BASILIC: an end-to-end pipeline that turns a compact `YAML` specification into a complete Bayesian analysis and a standardised, shareable report. BASILIC builds directly on `bilby` for likelihoods, samplers and compact binary approximants, and augments it with a curated registry of burst models, HTCondor-centred orchestration, deterministic injection tooling, and a coherent post-processing layer. The design choices were guided by a simple aim: to make model comparison routine rather than exceptional for short, information-rich signals.

Two use cases illustrate the scientific payoff. First, a multi-hypothesis re-analysis of GW231123 shows how a single configuration can test a state-of-the-art BBH model alongside several burst morphologies. In that application, the binary black hole interpretation is decisively favoured, while the power-law and cosmic string waveforms, though clearly preferred to noise, fail to match the time domain structure of the data. The exercise demonstrates the value of *contextual* model comparison: reporting a detection-versus-noise BF is not enough for classification when only a handful of cycles enter the band. Second, a controlled injection campaign quantifies recoverability and confusion across signal-to-noise ratio and morphology. Parameter posteriors stabilise once the network SNR approaches the typical detection threshold; BF scale as expected with SNR; and, crucially, there exist genuine ambiguity regimes in which very high-mass BBHs and low-index power-law bursts can earn comparable support from the same data. These studies underscore why tools that make cross-hypothesis checks trivial are needed *before* the interesting events arrive.

Along the way, several lessons emerged. (i) Prior choices can matter when data are short and structure-poor; BASILIC therefore surfaces priors explicitly and encourages side-by-side variants when a parameter (e.g. a spectral index) is weakly informed. (ii) Standardised outputs, such as per-model evidences and SNRs, comparable corner plots, and lightweight HTML summaries, all greatly shorten the path from samples to interpretation and collaboration. (iii) Reproducibility is not a luxury for campaigns: deterministic seeding, consistent PSD handling and fixed plotting conventions are essential to separate sampler idiosyncrasies from genuine data-driven trends.

The present implementation already supports the core workflows for short transients: single-event studies against multiple hypotheses and statistically robust injection campaigns. At the same time, the work highlights natural extensions. On the modelling side, broader burst coverage (additional CCSN-inspired templates, explicit memory variants, and simplified transients) and links to non-parametric reconstructions will widen the hypothesis set without over-fitting. On the inference side, trans-dimensional exploration (e.g. reversible-jump moves across a registry of models) would unify evidence ratios and parameter estimation within a single run, reducing orchestration overhead and enabling direct posterior weighting over model space. On the campaign side, automated summaries (efficiency curves, confusion matrices, and outlier triage) will turn large grids of runs into digestible, decision-making artefacts.

In short, BASILIC reframes burst inference from bespoke scripting to a reproducible, hypothesis-aware workflow that any GW scientist can operate: write one file, submit, and read a coherent report. The GW231123 application and the injection campaign demonstrate that this change in *process* unlocks better *science*: it makes it easy to ask the right comparative questions, to quantify uncertainty where morphology overlaps, and to scale from one interesting event to many. As detector sensitivity improves and short transients become more common and more diverse, the approach advocated here is not merely convenient; it is necessary for reliable source classification. The pipeline is available to the LVK, with source code hosted on GitLab and documentation that enables first runs without edits. With continued development along the lines outlined above, BASILIC can provide a practical foundation for systematic model comparison in the burst regime, and a bridge from candidate detection to robust, quantitative astrophysics.

A Basilic configuration for the GW231123 analysis

This appendix lists a self-contained YAML configuration that reproduces the GW231123 model comparison discussed in the main text. The file selects Hanford and Livingston cleaned strain via `gwpy`, computes a PSD from nearby off-source data, and evaluates five hypotheses: a binary black hole merger (`NRSur7Dq4`) and four burst-like alternatives (cosmic-string cusp, kink, kink-kink, and a generic power-law). Priors are loaded by BASILIC from the `priors/` directory using filenames that match the model names (e.g. `NRSur7Dq4_prior.yml`); they can be adjusted there as needed.

```
# ------------------------------
# Basilic configuration for GW231123
# ------------------------------

label: GW231123
output_path: ./GW231123_run

# Analysis band and sampling
detectors: [H1, L1]
duration: 10                # seconds
sampling_frequency: 4096   # Hz
detector_freq_band: [20, 2000]

# Sampler configuration
sampler: dynesty
request_cpus: 32
sampler_args:
  nlive: 2000
  dlogz: 0.1
  sample: rwalk
  walks: 20

# Hypotheses to test (must match available priors in priors/)
model_names:
```

```
  - NRSur7dq4
  - cosmic_string_cusp
  - cosmic_string_kink
  - cosmic_string_kink_kink
  - generic_power_law

# Data acquisition: pull calibrated, cleaned strain with gwpy
# Replace <TRIGGER_GPS> with the public trigger time (GPS seconds)
data_source:
  acquisition_method: gwpy
  channel_dict:
    H1: H1:GDS-CALIB_STRAIN_CLEAN_BAYESWAVE_S00
    L1: L1:GDS-CALIB_STRAIN_CLEAN_AR
  start_time: 1384782880.634
  end_time:   1384782890.634



# PSD strategy: read from dat (BayesWave psd)
psd_source:
  acquisition_method: dat
  path:
  H1: '/home/iuliudaniel.cuceu/misc/H1_PAS.dat'
  L1: '/home/iuliudaniel.cuceu/misc/L1_PAS.dat'

# HTCondor submission
request_cpus: 32
accounting_group: ligo.dev.o4.burst.allsky.stamp
conda_env_path: /home/iuliudaniel.cuceu/.conda/envs/basilic_env
```

`Duration` and `detector_freq_band` set the time-frequency window used in the likelihood. The `gwpy` acquisition pulls cleaned, calibrated strain for each IFO; The `psd_compute` block reads a PSD from a dat file; the parameters are typical for short-duration analyses and can be tuned. The list in `model_names` drives hypoth-

esis testing; BASILIC loads the corresponding priors from `priors/`. The sampler settings are conservative; the submission block informs HTCondor about the job it is supposed to run.

References

[1] J. Aasi et al., *Advanced LIGO*, *Class. Quant. Grav.* **32** (2015) 074001 [1411.4547].

[2] F. Acernese et al., *Advanced Virgo: a second-generation interferometric gravitational wave detector*, *Class. Quant. Grav.* **32** (2015) 024001 [1408.3978].

[3] T. Akutsu et al., *Overview of KAGRA: Detector design and construction history*, *PTEP* **2021** (2021) 05A101 [2005.05574].

[4] A.G. Abac et al., *GW231123: a Binary Black Hole Merger with Total Mass 190-265 $M_\odot$*, *arXiv* (2025) [2507.08219].

[5] I. Cuceu, M.A. Bizouard, N. Christensen and M. Sakellariadou, *GW231123: Binary Black Hole Merger or Cosmic String?*, *arXiv* (2025) [2507.20778].

[6] M. Maggiore, *Gravitational waves*, vol. 1, Oxford university press (2008).

[7] J. Skilling, *Nested Sampling*, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, R. Fischer, R. Preuss and U.V. Toussaint, eds., vol. 735 of *American Institute of Physics Conference Series*, pp. 395–405, Nov., 2004, DOI.

[8] J. Skilling, *Nested sampling for general Bayesian computation*, *Bayesian Analysis* **1** (2006) 833 .

[9] C. Cahillane and G. Mansell, *Review of the Advanced LIGO Gravitational Wave Observatories Leading to Observing Run Four*, *Galaxies* **10** (2022) 36 [2202.00847].

[10] A. Buikema et al., *Sensitivity and performance of the Advanced LIGO detectors in the third observing run*, *Phys. Rev. D* **102** (2020) 062003 [2008.01301].

[11] E. Capote et al., *Advanced LIGO detector performance in the fourth observing run*, *Phys. Rev. D* **111** (2025) 062002 [2411.14607].

[12] B.P. Abbott et al., *Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy*, *Phys. Rev. D* **93** (2016) 112004 [1604.00439].

[13] M. Maggiore, *Gravitational waves: volume 2: astrophysics and cosmology*, Oxford University Press (2018).

[14] C. Palenzuela, *Introduction to Numerical Relativity*, *Front. Astron. Space Sci.* **7** (2020) 58 [2008.12931].

[15] F. Pretorius, *Evolution of binary black hole spacetimes*, *Phys. Rev. Lett.* **95** (2005) 121101 [gr-qc/0507014].

[16] M. Campanelli, C.O. Lousto, P. Marronetti and Y. Zlochower, *Accurate evolutions of orbiting black-hole binaries without excision*, *Phys. Rev. Lett.* **96** (2006) 111101 [gr-qc/0511048].

[17] J.G. Baker, J. Centrella, D.-I. Choi, M. Koppitz and J. van Meter, *Gravitational wave extraction from an inspiraling configuration of merging black holes*, *Phys. Rev. Lett.* **96** (2006) 111102 [gr-qc/0511103].

[18] S. Husa, S. Khan, M. Hannam, M. Pürrer, F. Ohme, X. Jiménez Forteza et al., *Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal*, *Phys. Rev. D* **93** (2016) 044006 [1508.07250].

[19] S. Khan, S. Husa, M. Hannam, F. Ohme, M. Pürrer, X. Jiménez Forteza et al., *Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era*, *Phys. Rev. D* **93** (2016) 044007 [1508.07253].

[20] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme et al., *Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms*, *Phys. Rev. Lett.* **113** (2014) 151101 [1308.3271].

[21] A. Taracchini et al., *Effective-one-body model for black-hole binaries with generic mass ratios and spins*, *Phys. Rev. D* **89** (2014) 061502 [1311.2544].

[22] A. Nagar et al., *Time-domain effective-one-body gravitational waveforms for coalescing compact binaries with nonprecessing spins, tides and self-spin effects*, *Phys. Rev. D* **98** (2018) 104052 [1806.01772].

[23] M. Breschi, S. Bernuzzi, F. Zappa, M. Agathos, A. Perego, D. Radice et al., *kiloHertz gravitational waves from binary neutron star remnants: time-domain model and constraints on extreme matter*, *Phys. Rev. D* **100** (2019) 104029 [1908.11418].

[24] R. Jeannerot, J. Rocher and M. Sakellariadou, *How generic is cosmic string formation in SUSY GUTs*, *Phys. Rev. D* **68** (2003) 103514 [hep-ph/0308134].

[25] B.P. Abbott et al., *Constraints on cosmic strings using data from the first Advanced LIGO observing run*, *Phys. Rev. D* **97** (2018) 102002 [1712.01168].

[26] R. Abbott et al., *Constraints on Cosmic Strings Using Data from the Third Advanced LIGO–Virgo Observing Run*, *Phys. Rev. Lett.* **126** (2021) 241102 [2101.12248].

[27] K.G. Arun et al., *New horizons for fundamental physics with LISA*, *Living Rev. Rel.* **25** (2022) 4 [2205.01597].

[28] A.K. Divakarla, E. Thrane, P.D. Lasky and B.F. Whiting, *Memory Effect or Cosmic String? Classifying Gravitational-Wave Bursts with Bayesian Inference*, *Phys. Rev. D* **102** (2020) 023010 [1911.07998].

[29] D. Christodoulou, *Nonlinear nature of gravitation and gravitational wave experiments*, *Phys. Rev. Lett.* **67** (1991) 1486.

[30] K.S. Thorne, *Gravitational-wave bursts with memory: The Christodoulou effect*, *Physical Review D* **45** (1992) 520.

[31] Y.B. Zel'dovich and A.G. Polnarev, *Radiation of gravitational waves by a cluster of superdense stars*, *Soviet Astronomy* **18** (1974) 17.

[32] I. Cuceu and S. Robles, *Direct nonparametric multimessenger constraints on the equation of state of cold dense nuclear matter*, *Physical Review D* **111** (2025) 123029.

[33] L. Blanchet and T. Damour, *Hereditary effects in gravitational radiation*, *Phys. Rev. D* **46** (1992) 4304.

[34] C. Talbot, E. Thrane, P.D. Lasky and F. Lin, *Gravitational-wave memory: Waveforms and phenomenology*, *Phys. Rev. D* **98** (2018) 064031.

[35] C.J. Richardson, M. Zanolin, H. Andresen, M.J. Szczepańczyk, K. Gill and A. Wongwathanarat, *Modeling core-collapse supernovae gravitational-wave memory in laser interferometric data*, *Phys. Rev. D* **105** (2022) 103008 [2109.01582].

[36] C. Afle and D.A. Brown, *Inferring physical properties of stellar collapse by third-generation gravitational-wave detectors*, *Phys. Rev. D* **103** (2021) 023005 [2010.00719].

[37] E. Abdikamalov, S. Gossan, A.M. DeMaio and C.D. Ott, *Measuring the Angular Momentum Distribution in Core-Collapse Supernova Progenitors with Gravitational Waves*, *Phys. Rev. D* **90** (2014) 044001 [1311.3678].

[38] J. Powell and B. Müller, *Inferring astrophysical parameters of core-collapse supernovae from their gravitational-wave emission*, *Phys. Rev. D* **105** (2022) 063018 [2201.01397].

[39] L.O. Villegas, C. Moreno, M.A. Pajkos, M. Zanolin and J.M. Antelis, *Parameter estimation from the core-bounce phase of rotating core collapse supernovae in real interferometer noise*, *Class. Quant. Grav.* **42** (2025) 115001 [2304.01267].

[40] T. Eccleston and M.C. Edwards, *Generative adversarial network for stellar core-collapse gravitational waves*, *Phys. Rev. D* **110** (2024) 104055 [2408.02895].

[41] LIGO Scientific Collaboration, "LIGO Algorithm Library - LALSuite." free software (GPL), 2018. 10.7935/GT1W-FZ16.

[42] K. Wette, *SWIGLAL: Python and Octave interfaces to the LALSuite gravitational-wave data analysis libraries*, *SoftwareX* **12** (2020) 100634.

[43] B.P. Abbott et al., *All-sky search for short gravitational-wave bursts in the third advanced ligo and advanced virgo run*, *Phys. Rev. D* **104** (2021) 122004 [2107.03701].

[44] B.P. Abbott et al., *All-sky search for short-duration gravitational-wave bursts in the first part of the fourth ligo–virgo–kagra observing run*, 2507.12374.

[45] G. Ashton, M. Hübner, P.D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu et al., *BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy*, *Astrophys. J. Suppl.* **241** (2019) 27 [1811.02042].

[46] I.M. Romero-Shaw, C. Talbot, S. Biscoveanu, V. D'Emilio, G. Ashton, C.P.L. Berry et al., *Bayesian inference for compact binary coalescences with BILBY: validation and application to the first LIGO-Virgo gravitational-wave transient catalogue*, *Mon. Not. Roy. Astron. Soc.* **499** (2020) 3295 [2006.00714].

[47] M. Drago, S. Klimenko, C. Lazzaro, E. Milotti, G. Mitselmakher, V. Necula et al., *Coherent waveburst, a pipeline for unmodeled gravitational-wave data analysis*, *SoftwareX* **14** (2021) 100678.

[48] S. Klimenko, I. Yakushin, A. Mercer and G. Mitselmakher, *Coherent method for detection of gravitational wave bursts*, *Class. Quant. Grav.* **25** (2008) 114029 [0802.3232].

[49] S. Klimenko et al., *Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors*, *Phys. Rev. D* **93** (2016) 042004 [1511.05999].

[50] T. Bayes, *An essay towards solving a problem in the doctrine of chances*, *Philosophical Transactions of the Royal Society of London* **53** (1763) 370.

[51] N. Christensen and R. Meyer, *Parameter estimation with gravitational waves*, *Rev. Mod. Phys.* **94** (2022) 025001 [2204.04449].

[52] H. Jeffreys, *The theory of probability*, OuP Oxford (1998).

[53] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller, *Equation of state calculations by fast computing machines*, *J. Chem. Phys.* **21** (1953) 1087.

[54] W.K. Hastings, *Monte carlo sampling methods using markov chains and their applications*, *Biometrika* **57** (1970) 97.

[55] R.M. Neal, *Mcmc using hamiltonian dynamics*, in *Handbook of Markov Chain Monte Carlo*, S. Brooks, A. Gelman, G. Jones and X.-L. Meng, eds., pp. 113–162, Chapman and Hall/CRC (2011).

[56] J. Skilling, *Nested sampling for general bayesian computation*, *Bayesian Analysis* **1** (2006) 833.

[57] J.S. Speagle, *dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences*, *Mon. Not. Roy. Astron. Soc.* **493** (2020) 3132 [1904.02180].

[58] I. Cuceu, "Basilic: Bayesian transient modeling and classification."

[59] R. Abbott et al., *GW190521: A Binary Black Hole Merger with a Total Mass of* $150M_\odot$, *Phys. Rev. Lett.* **125** (2020) 101102 [2009.01075].

[60] R. Abbott et al., *Properties and Astrophysical Implications of the 150* $M_\odot$ *Binary Black Hole Merger GW190521*, *Astrophys. J. Lett.* **900** (2020) L13 [2009.01190].

[61] V. Varma et al., *Surrogate models for precessing binary black hole simulations with unequal masses*, *Phys. Rev. Research.* **1** (2019) 033015 [1905.09300].

[62] LVK, "GWOSC DATA GW231123, doi.org/10.7935/anj7-6q40."

Acknowledgements

The author gratefully acknowledges Marie-Anne Bizouard, Nelson Christensen, Mairi Sakellariadou and Viviana Fafone, without whom much of this would not have been possible. Furthermore, the author acknowledges support through an EMJM scholarship funded by the European Union in the framework of the Erasmus+, Erasmus Mundus Joint Master in Astrophysics and Space Science – MASS. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or granting authority European Education and Culture Executive Agency (EACEA). Neither the European Union nor the granting authority can be held responsible for them.