



1506
UNIVERSITÀ
DEGLI STUDI
DI URBINO
CARLO BO

UNIVERSITÀ DEGLI STUDI DI URBINO CARLO BO

Department of Pure and Applied Sciences

Ph.D. PROGRAMME IN: RESEARCH METHODS IN SCIENCE AND TECHNOLOGY

CYCLE XXXVIII

**Advancing gravitational-wave searches from compact binary
coalescences through novel algorithms and machine learning.**

ACADEMIC DISCIPLINE: PHYS-01/A

Coordinator: Prof. Luca Lanci

Supervisor: Prof. Gianluca Maria Guidi

Ph.D. student: Lorenzo Mobilia

ACADEMIC YEAR
2024/2025

Abstract

The first detection of Gravitational Wave event GW150914, by LIGO-Virgo-Kagra (LVK) collaboration opened the era of gravitational-wave astronomy. This discovery, occurring a century after Einstein’s prediction of gravitational waves in General Relativity, determined a scientific breakthrough. Since then, 218 events have been reported up to the first half of the current observing run (O4). These observations have enabled significant scientific advancements, including the first binary-neutron-star detection GW170817, the candidate intermediate-mass black-hole merger GW190521, and the most massive binary event to date, GW231123. Achieving such measures depends on continuous hardware upgrades and methodological improvements across the LVK detectors and data-analysis pipelines. Key challenges in detecting such extremely faint signals include noise sources characterization and isolation and developing algorithms that can distinguish transient noise from genuine astrophysical signals.

In this context, we focus on data analysis for compact binary coalescences, using the Multi-Band Template Analysis (MBTA) pipeline, which employs matched-filtering to trigger signals in the interferometer data streams and it is currently deployed in observing runs. We investigate alternative strategies for MBTA template-bank construction that substantially reduce bank-generation time while maintaining high signal-recovery efficiency. We also assess a machine learning approach in which a Random Forests (RF) classifier is applied to MBTA triggers to separate noise from signal more effectively. Using O4a single-detector triggers, we compare RF-based statistics with respect to the MBTA ranking-statistics. We then estimate the astrophysical source probability p_{astro} from the RF-derived statistics for coincident MBTA triggers in O3 and compare these values with the probabilities reported in the catalogues.

Finally, we explore proof-of-concept Deep Learning methods for Gravitational Waves detection. We construct feature-rich representations from matched-filtering output statistics that can be learned by a Convolutional Neural Network, yielding an alternative triggering approach. This method offers advantages relative to the state-of-the-art pipelines: it requires a much smaller template bank and does not rely on the χ^2 consistency test to reject noise candidate triggers.

Introduction

τοῦτο γάρ ἐστιν ὁ χρόνος, ἀριθμὸς κινήσεως κατὰ τὸ πρότερον καὶ ὕστερον.

“This is the time, the number of movement with respect to before and after.”

ARISTOTLE, *Physics* IV, 11, 219b1–2.

Scientific progress has advanced alongside humanity’s drive to understand nature. From the twentieth-century discoveries of Quantum Mechanics — describing the imperceptibly small, such as electrons and atoms — to Einstein’s theory of gravitation, which governs the immensely vast cosmic structures like galaxies and the space-time itself, questions about Universe has been a constant in our history. Einstein’s description of the space-time, known as General Relativity, provides a framework for the universe’s evolution and for many astrophysical phenomena. In particular It predicts travelling ripples in space-time — the Gravitational Waves — emitted by accelerating compact objects such as neutron stars and black holes, which were directly detected a century later, marking a major scientific breakthrough.

Achieving such discovery required decades of scientific and technological development: sustained efforts, repeated trials — many of them unsuccessful — and struggle of thousands of scientists. Only in recent years have laser interferometers become sensitive enough to detect such faint signals. In 2015 the first Gravitational Wave was observed, and since then the progress has accelerated: observations of binary black-holes and binary-neutron-star mergers have clarified the nature of extreme compact objects and helped explain the origin of certain heavy elements. Cosmological parameters, including the Hubble constant, are now being probed with these data. Among the various astrophysical sources of Gravitational Waves, compact binary coalescences — merging pairs of black-holes or neutron-stars — are especially important: their relatively high rates and detectability with current instruments make them prime laboratories for precision tests of fundamental physics. These events are the primary focus of this work.

Detecting such signals, catastrophic in their intrinsic nature, demands extraordinary advances in both technology and statistical methodology. Alongside the hardware push to build ever more sensitive interferometers, algorithms and software for gravitational-wave detection play a crucial role. In particular, machine learning and Artificial Intelligence are increasingly integrated into current searches. These techniques also occupy an important place in this study.

This thesis advances gravitational-wave searches by enhancing and developing methods for detecting of compact binary coalescences. The work was conducted during the fourth observing run

(O4) of the LIGO–Virgo–KAGRA Collaboration, of which I am on current date an active member. It investigates strategies to optimize a key component of the compact binary coalescence searches: the template bank for the currently deployed Multi-Band Template Analysis (MBTA) pipeline. This requires fast algorithms for large dataset in non-Euclidean parameter spaces. In parallel, the thesis explores approaches to improve detection performance, including machine learning enhancements to MBTA and Deep Learning models for direct signal identification. The manuscript is organized as follows:

- Chapter 1 provides a concise introduction to Gravitational Waves both historically and theoretically: it sketches a derivation from the Einstein field equations and interprets the resulting solutions as measurable physical phenomena. It also surveys expected compact binary coalescence sources and their orders of magnitude.
- Chapter 2 reviews the currently operating laser-interferometric detectors and the principal noise sources that affect them. It then introduces the matched-filtering procedure required to extract Gravitational Waves signals and describes the MBTA pipeline.
- Chapter 3 formulates the template bank construction problem and presents MBTA-specific bank architecture, together with the associated algorithmic developments and analytical studies.
- Chapter 4 addresses machine learning approaches to enhance the detection efficiency of the MBTA pipeline, using single-detector triggers from the first part of O4 and two-detector coincidences from O3. Furthermore, it investigates the estimation of the probability of astrophysical origin for candidates and compares the results with official values reported in catalogues.
- Chapter 5 explores a Deep Learning approach for the direct detection of gravitational-wave signals. Building on the matched-filtering framework with a reduced-size template bank, it reports an extensive simulation campaign and a comparison against classical search methods.

Lorenzo Mobilia
December 2025

Contents

List of Figures	ix
List of Tables	xi
1 Gravitational Waves: discovery and theory	1
1.1 Gravitational-wave science: a brief history	1
1.2 Linearization of Einstein Equations	5
1.3 The generation mechanism in linearized theory	12
1.4 The inspiral of compact objects	15
1.5 Main observational results	21
1.6 Sources and detection	22
2 Detection of the inspiral binaries: instruments, techniques and pipelines	25
2.1 Interferometers	25
2.2 Noise sources	28
2.2.1 Shot noise	29
2.2.2 Radiation pressure	30
2.2.3 Other noise sources	31
2.3 Detecting the Gravitational Waves - Data analysis and techniques	33
2.3.1 Matched-filtering	35
2.3.2 The matched-filtering statistics	36
2.3.3 Coalescence of compact binaries	38
2.3.4 Mismatch	40
2.4 The Multi-Band Template Analysis Pipeline	42
2.4.1 Single triggers and coincidence	44
2.4.2 False Alarm Rate and probability of being astrophysical source	46
3 Template bank construction	49
3.1 Geometric, stochastic and hybrid placement	50
3.2 Real Template bank production optimization	52
3.2.1 Algorithm design and implementation	52

3.2.2	First study: the BNS Virtual Template case	55
3.3	The analytical coverage	62
3.3.1	Relationship between the match and the frequency bands	63
3.3.2	Properties of α and β	65
3.3.3	Match sharing among the frequency bands	66
3.3.4	Coverage problem among different frequency bands	67
3.4	The UBER BANK case	69
3.4.1	Impact of the f_{cut} choice	71
4	Random Forests for MBTA	79
4.1	Machine learning for CBC detection	80
4.2	Random Forests algorithm	81
4.2.1	Tree based methods	82
4.2.2	Bootstrap and bagging	83
4.2.3	Random Forests	84
4.3	Random Forests for MBTA - single triggers	87
4.3.1	Dataset, features and hyper-parameters	88
4.3.2	Single triggers results	90
4.4	Random Forests for MBTA - coincidence triggers	97
4.5	Random Forests for MBTA - p_{astro} computation	101
5	Convolutional Neural Network for gravitational-wave triggering	109
5.1	Convolutional Neural Network	110
5.2	The Time-Template map	115
5.3	ResNet, training and results	123
5.3.1	Simulation 1	125
5.3.2	Simulation 2	128
5.3.3	Simulation 3	131
6	Conclusions	135
	Bibliography	139

List of Figures

1.1	First Detection of GW150914 [1] observed by LIGO Hanford (left panels) and LIGO Livingston (right panels). The first row shows the strain time series band-pass filtered in the 35 – 350Hz. The second row displays waveform reconstruction obtained from parameters consistent with the event, numerical–relativity and wavelet-based reconstructions [2–4]. The third row shows the residual obtained after subtracting the numerical-relativity waveform from the data. The bottom row presents a time–frequency representation of the data, highlighting the characteristic chirp behavior of the signal [5].	4
1.2	(Top) Observing plans and operational/planned instrument : The first four observing runs have been completed, yielding a total of 218 recorded events up to the first half of O4 (O4a). Looking ahead, O5 is expected to begin in mid-2027 and continue until approximately mid-2030. (Bottom) IFOs network ground based detectors operational/planned: The figure summarizes the current global network of laser-interferometer observatories, including existing facilities as well as proposed upgrades and next-generation instruments that will extend the astrophysical reach and sensitivity of gravitational-wave astronomy.	5
1.3	Motion of test particles in + and x polarization polarization	11
1.4	Geometry of $\vec{x}, \vec{y} \in \vec{x} - \vec{y}$ inside the body of typical length L	12
1.5	Different source types, assuming the same starting frequency, produces signals with different durations: hundreds of seconds for a BNS system, and few tens of seconds for a BBH. In this example the merger time is $t_c = 40s$. The spectrogram also highlights the different chirp behavior of the two systems.	20
2.1	Example of Michelson-Morlet IFO - credits to [6]	26
2.2	Example of MM IFO with FP cavity - credits to [6]	28
2.3	The scheme of the superattenuator. A Virgo mirror is supported in a vacuum by this kind of structure, which dampens seismic vibrations. It is a chain of pendula hanging from an upper platform and supported by three legs attached to ground, forming an inverted pendulum. — credits to [7]	32
2.4	Comparison of quantum noise (Top) and ALIGO noise budget (Bottom).	33

2.5	39
2.6	Matched-Filtering time series for an injected BBH of $5 - 5M_{\odot}$ at $800Mpc$ in Gaussian noise colored by the PSD of LIGO IFO. For simulations we use the Advanced LIGO zero-detuned high-power design sensitivity curve (aLIGOZeroDetHighPower) as implemented in LALSimulation [8].	41
3.1	Flowchart of the algorithm to construct the RT bank by retaining only non-redundant templates. N is the number of templates. The set Ignored contains indices of redundant templates; these are associated to a real template and excluded from subsequent iterations.	53
3.2	Match distributions for a binary neutron star template bank batch as a function of the chirp-mass difference ΔM_c with respect to the reference template T_0 . Panels (a) and (b) show the distributions of matches computed in the low-frequency (20–80 Hz) and high-frequency (80–1024 Hz) bands, respectively, for 15,000 non-spinning templates. Panels (c) and (d) display the match as a function of both ΔM_c and the mass ratio q , with the colour scale indicating the match value. In the low-frequency band, the match is primarily governed by ΔM_c and remains nearly constant across different values of q at fixed chirp mass. In contrast, in the high-frequency band the match exhibits a stronger dependence on q , leading to a broader distribution near $\Delta M_c \simeq 0$. Only a small fraction of templates achieve matches above 0.97, concentrated close to the chirp mass of T_0 . Templates are ordered in ascending chirp mass, hence $\Delta M_c < 0$.	54
3.3	Schematic representation of the map–reduce procedure applied to the VT bank to compute the Real Template (RT) banks. Starting from the Virtual Template Bank, we first divide it into N batches (map phase). In each batch, the matches among templates within the considered frequency bandwidth are computed, and when the match exceeds a given threshold, the corresponding template is discarded and its index recorded (reduce phase). Finally, the reduced batches are merged to obtain the final RT bank. For reference, only the first template is highlighted for each batch from which the procedure starts.	55
3.4	Virtual Template bank for the BNS case.	56
3.5	f_{cut} distribution for the BNS bank. The median value — highlighted in red — indicated the frequency cut chosen for building the RT banks.	57
3.6	Example of procedure for selecting the frequency cut for a single template. In this case a BBH ($m_1 = 46.81M_{\odot} - m_2 = 4.26M_{\odot} - \chi_1^z = \chi_2^z = -0.803$) is taken into account. The blue curve represent the cumulative SNR ² (σ^2), the red line is the 50% – SNR ² point and the green-dot line is the corresponding frequency.	57
3.7	Map-reduce procedure applied to the VT bank for computing the RT banks result.	59
3.8	Bank Vericator procedure applied to the RTLf and RTHF banks (a) and the BNS VT bank (b)	60

3.9	Scatter plot for injected masses. The colorbar represents the match gradient for each injection. We see that no particular behavior emerges from the distribution, the injections are uniformly recovered within the expected fitting factor.	61
3.10	Distribution of f_{cut} among the UBER BANK. The red line identifies the median frequency cut and correspond to the chosen frequency for bandwidth selection. . . .	71
3.11	Map-Reduce algorithm results for the VT UBER BANK to obtain the RT Banks. . . .	72
3.12	Cumulative fitting factor obtained for the RT banks and the UBER BANK in the whole parameter space of a full search.	73
3.13	Scatter plot for injected masses. The colorbar represents the match gradient for each injection. Even though the coverage is within the expected values, we see that the degradation of the fitting factor associated to the high-mass parameter space of BBH in the LF band.	74
3.14	Scatter plot of injected masses with Fitting Factor (Match) for the full-bandwidth Virtual Template bank (Focus on BBH parameter space). We see that at top right corner the VT bank fails to recover properly the injected signals, this means that the RT banks that are constructed from this one will inherit the same coverage behavior. . . .	75
3.15	Distribution of f_{cut} among the UBER BANK in the $q - \chi_{\text{eff}}$ space. The colorbar shows a gradient with descending frequency values toward high q and low χ_{eff} values. . . .	76
3.16	Pictorial representation of the impact of q for T_0 on the SNR distribution among the frequency bands. On top (a) the $f_{\text{cut}} = 70\text{Hz}$ shows how the SNR distribution is unbalanced in favor of the LF band while in (b) the frequency cut imposed is $f_{\text{cut}} = 43\text{Hz}$ with an SNR equally shared among the bands. The waveforms are obtained using SEOBNRv4_ROM in the frequency interval specified for the SNR study. . . .	77
3.17	Top: scatter plot with UBER BANK masses and corresponding percentage of SNR percentage for each template (LF - left; HF - right). Bottom: distribution of SNR percentage for the two frequency bands. The analysis is performed using the nominal frequency cut of 80 Hz adopted by MBTA.	78
4.1	Left image shows a partition of two-dimensional feature space by recursive binary splitting, as used in CART. Right image shows the corresponding tree. Credits to [9]	82
4.2	Schematic of the bootstrap process. Datasets Z^* have all the same size N and are drawn from replacement from Z	84
4.3	Example of tree in a Random Forests used for this study up to three level-depth decision nodes.	85
4.4	Features pairplot distribution from a subset of 5000 events from dataset. We expect that the Random Forests algorithm is capable in combining these informations of the different population and constructing a statistics that separates the injections from the noise labeled triggers.	90

4.5	Distributions of ranking-statistics (amplitude) and p_s for L and H detectors. The top plots show ranking-statistics distributions; the bottom plots show p_s distributions for the same test dataset.	93
4.6	Receiver Operating Characteristic (ROC) curves for the best models of O4a.	94
4.7	Efficiency in injection recovery ratio between p_s and ρ_{rw} for $F2$ features set. Events binned in M_{tot} and χ_{eff} . Each box shows $\log_{10}(R)$, where $R = \frac{N_d(\tilde{p}_s)}{N_d(\tilde{\rho}_{rw})} \Big _{\tilde{\alpha}}$. Green gradient indicates $R > 0$ (higher recovery with p_s), conversely red gradient is for $R < 0$ (higher recovery with ρ_{rw}) and white bins are where no injections at all were performed.	95
4.8	Injection and noise population for H and L. The net density of injection of the upper left corner of the plot with respect to the bottom may give some hints about the lack efficiency in detection for H, yet we do not appreciate a significant difference between the IFOs regarding the population distribution, this may suggest that other causes affects this H mis-recovery.	96
4.9	Distributions of ranking-statistics (ρ_{rw}) and p_s for O3a and O3b HL-Von coincidence MBTA triggers. The top plots show ranking-statistics distributions; the bottom plots show p_s distributions for the same test dataset.	99
4.10	Receiver Operating Characteristic (ROC) curves for HL-Von coincidence triggers (O3). Even though the red curve obtained from the p_s statistics is almost everywhere above the ranking-statistics one, the difference is not substantial as we saw in the single detector case.	100
4.11	Probability density function obtained for O3a \tilde{p}_s distributions for signal (red) and (blue) labeled triggers.	103
4.12	p_{astro} values obtained via Kernal Density Estimation over test dataset considering the \tilde{p}_s statistics for O3a (left) and O3b (right) triggers. We see that all the noise-labelled signal lay in the bottom part of the plot, with no triggers above 0.15.	104
4.13	Comparison between $p_{astro}^{(p_s)}$ values obtained by p_s statistics and p_{astro} for O3 MBTA catalogues' events.	105
4.14	$p_{astro}^{(p_s)}$ and p_{astro} without ER feature - O3a MBTA catalogue's events	106
5.1	Illustrations of (a) the perceptron structure and (b) the mathematical formalism behind neural network layers.	110
5.2	Example of a Neural Network with 5 layers. The green light dots represent the input, the light violet are the hidden layers and the light red one is the output layer.	111
5.3	Illustrations of a convolutional operation (a) and the result of a series of kernel applied to an input image (b).	114
5.4	Illustrations of (a) the Convolutional Neural Network structure and (b) the shortcut connection that characterize the Residual Block	114

5.5	Effect of the χ^2 test on $\rho(t)$ for different template–data configurations. (a) Perfect match, (b) template matched to a sine-Gaussian glitch, and (c) template-signal mismatch. The reweighed statistic ρ_{rw} penalizes non-signal-like transients, significantly reducing their significance.	117
5.6	Spectrograms showing template–template correlations. We used a reference BNS with $M_c^{\text{ref}} = 2.17M_\odot$ and compare systems with chirp-mass below (top) and above (bottom) the reference. Each panel shows the time–frequency representation of the reference signal (central track) together with two realizations of the same template chirp mass, shifted in time with respect to the reference coalescence time t_0 . The presence of three chirp tracks therefore reflects the reference signal and two time-shifted instances of the template waveform, used to illustrate how the relative merger time determines whether the template intersects the reference chirp. The reference chirp track is intersected by a lower M_c system only if it merges after the reference time t_0 ; conversely, a higher M_c system crosses the reference track if it merges before t_0 . 120	
5.7	Examples of the TT-SNR-Map obtained by piling up the $\rho(t)$. On Top the picture corresponding to Gaussian Noise solely, on the center an injected signal and on the bottom an injected signal with a glitch superimposed. The horizontal and vertical axis represent, respectively, the time interval around the injection time and the chirp-mass of the template used to matched-filter the data.	121
5.8	Architecture of the Residual Network used in this work (Top) and focus on the Residual Block (Bottom).	124
5.9	Results for purely Gaussian Noise and BNS injections.	126
5.10	BNS signals TT-Map injected, respectively, at 50 (a), 100 (b), 200 (c) and 400 (d) Mpc. We see that the farther the signal is, the fader the TT-Map will result, consequently really far away signals cannot be distinguished from pure Gaussian noise.	127
5.11	Simulation 1: ROC curves for r (CNN-output) and ρ^{max} statistics.	128
5.12	Injected BNS signal in Gaussian noise with sine-Gaussian glitch (top) and zoom over glitch (bottom)	130
5.13	Simulation 2: ROC curves for $\rho_{\text{rw}}^{\text{max}}$, r (CNN-output) and ρ^{max} statistics.	131
5.14	TT-Map constructed considering $\rho(t)$ (top), $\chi_r^2(t)$ (center) and $\rho_{\text{rw}}(t)$ (bottom) statistics. The horizontal and vertical axes are the same as in Fig. 5.7.	133
5.15	Simulation 3: ROC curves for $\rho_{\text{rw}}^{\text{max}}$, r (CNN-output) and ρ^{max} statistics.	134

List of Tables

3.1	BNS production of Virtual Template Bank parameters - ASD reference: [8].	56
3.2	Comparison of Virtual Template Bank Parameters	70
4.1	MBTA triggers - O4a	88
4.2	hyper-parameters grid search - O4a	91
4.3	Best Models - O4a	91
4.4	Features sets	92
4.5	O3 datasets	97
4.6	hyper-parameters grid search- O3	98
4.7	Best Models - O3	98
4.8	O3a Events p_{astro} comparison	107
4.9	O3b Events p_{astro} comparison	108
5.1	Parameter ranges used for the injected sine-Gaussian signals.	129

Chapter 1

Gravitational Waves: discovery and theory

1.1 Gravitational-wave science: a brief history

On the 14th of September 2015, for the first time in history, a gravitational-wave (GW) signal, GW150914, was detected by the GW detector network composed by the two LIGO interferometers (IFOs), both located in the United States of America [1] (Fig. 1.1). This measurement is a milestone in the history of physics. A further confirmation of Einstein's theory of gravity, General Relativity (GR), that took over one century to find direct experimental confirmation. The detection of this event marked the birth of gravitational-wave astronomy, allowing mankind a deeper comprehension of the universe. In the following years, other capital events have been recorded, as the first Binary Neutron Star (BNS) merging system GW170817 [10]. Such event empowered multi-messenger astrophysics, combining information from the GW and the electro-magnetic counterpart. Another honorable mention is the first intermediate-mass black hole detection, resulting from a merger of two black holes that generated GW190521 [11]. Today, 218 events have been recorded [12–16], and huge scientific advancements have been pushed forward.

These exciting achievements have been possible thanks to the outstanding advancement in technology that permitted the interferometers to access a sufficient level of sensibility, which includes noise isolation, engineering developments, and decades of research. But the history of GWs has not been so linear and easy. Before the General Relativity theory, developed by Einstein in 1915, there were several speculations around the possible existence of waves due to gravity. Most of those thoughts were due to the formalization of the Electro Magnetic field by James Clerk Maxwell, the Scottish physicist, who predicted the existence of Electro Magnetic waves [17]. The use of the conceptual and formal tool of fields was a step in the right direction to understand gravity, since the assumption of action at a distance developed by Newton in the 17th century was no longer sufficient. The first hypothesis of GWs existence, following this path, was due to H. Poincaré [18] in 1905, but to have a

consistent and solid comprehension of this phenomenon, we have to wait until the General Relativity theory of Einstein [19] in its present form.

The GWs, predicted by Einstein's theory of General Relativity, were initially thought by the scientist himself to be merely a mathematical artifact. It is worth mentioning the mail exchanged with Schwarzschild in 1916, where Einstein asserted the non-existence of GWs, or, in equivalent terms, 'non-existence of a gravitational dipole'. In any case, the General Relativity theory was indeed too strong for Einstein's doubts too, and the integration of its equations still showed solutions compatible with the production of GWs. The theoretical debate continued until 1937, when Einstein published the article *On Gravitational Waves* [20], but even after this publication, the father of the General Relativity theory, was quite skeptical as shown by his conclusion during a 1936 Princeton lecture: "If you ask me whether there are gravitational waves or not, I must answer that I don't know. But it is a highly interesting problem."

Only decades later the idea of measuring these events did take form. Along with other doubts, in fact, a huge debate concerned the possibility that a GW could transport energy—a problem that found consensus in the scientific community only in the '60s. These problems were discussed at GR1, the famous meeting at Chapel Hill in 1957. At the conference some physicists joined like R. Feynman and J. Weber. The former, in order to convince the audience of the possibility for a GW to transmit energy, formulated the 'sticky bead' argument, that is a thought experiment designed to show that gravitational radiation is predicted by General Relativity and it has indeed physical effects. Consequently, Weber started developing some ideas for measuring the gravitational waves, such as the resonant bar detectors [21–23]. Unfortunately, those instruments, in principle capable of capturing such signals, could not access the needed sensitivity to effectively measure astrophysical events.

To obtain the first indirect measurement of Gravitational Waves, the Hulse–Taylor binary pulsar PRS B1913+16 measurement was needed in the late '70s. This pioneering measure consisted in observing the variation of the arriving time of the pulses, both in delaying and anticipating signals. Such an effect is consistent with the variation of the orbit due to the loss of energy caused by the emission of Gravitational Waves. [24]. To directly measure GWs, an interferometer is needed. Weber himself wrote a note about the possibility of using such instruments, even though he explored practically only the resonant bar detectors. The first prototypes of IFOs were produced from the late '60s, and after decades of constant amelioration, this technology became sensitive enough to measure the physical scale needed to detect GWs.

Right now, there is a network of IFOs running in data acquisition campaign, the two LIGO [25] in the US, Virgo [26] in Italy and GEO600 in Germany [27], even though the latter is used mostly as benchmark for technology test rather than effective GW measurements. In addition, an underground interferometers is acquiring data in Japan (KAGRA [28]) and another LIGO is foreseen to be built in India [29], see Fig. 1.2b for reference. Several observing run, with progressively improved detectors' sensitivity, have been conducted and the current observation plan is shown in Fig. 1.2a. In the future another, generation of IFOs will be available, the Einstein Telescope in Europe [30] and the CosmicExplorer in the US [31, 32]. Those instruments will access lower frequencies and allow us to

measure an incredible amount of new possible sources, with new outstanding scientific discoveries. To conclude this panoramic, it also necessary to mention LISA [33], a space-array interferometer, which trough accessing extremely low frequency bands, will open our eyes to new promising physics.

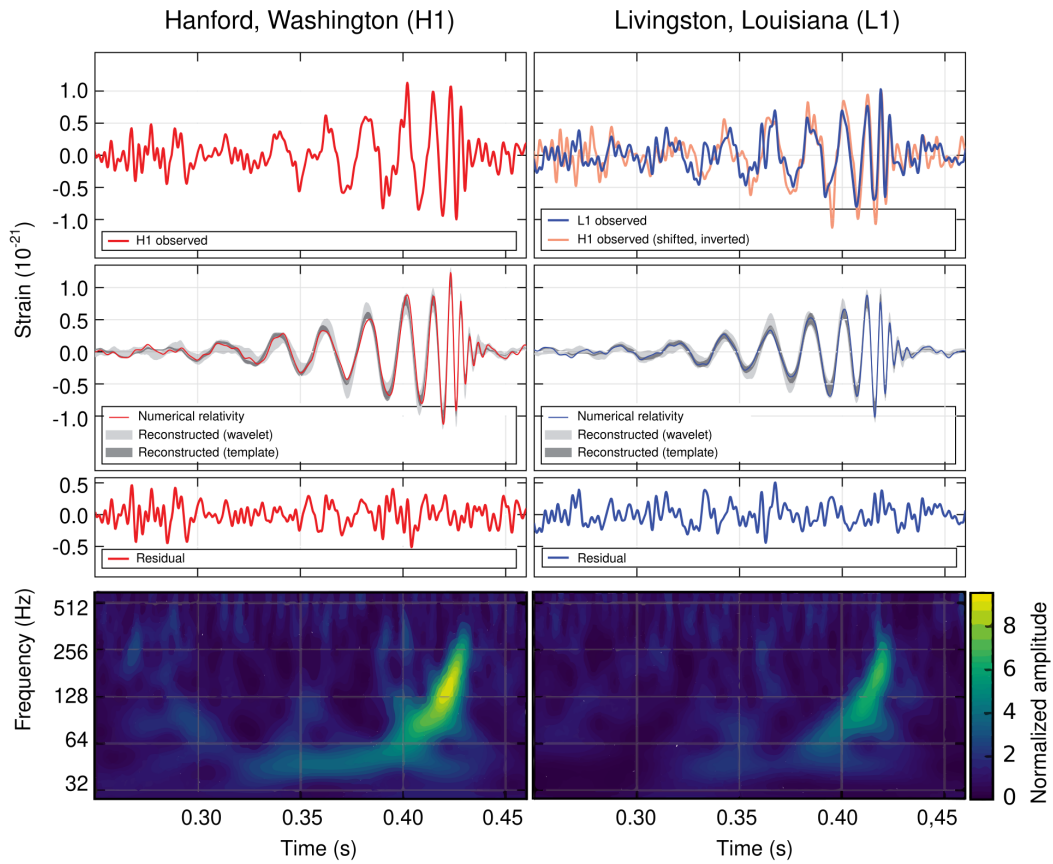




Figure 1.1 First Detection of GW150914 [1] observed by LIGO Hanford (left panels) and LIGO Livingston (right panels). The first row shows the strain time series band-pass filtered in the 35 – 350Hz. The second row displays waveform reconstruction obtained from parameters consistent with the event, numerical–relativity and wavelet-based reconstructions [2–4]. The third row shows the residual obtained after subtracting the numerical-relativity waveform from the data. The bottom row presents a time–frequency representation of the data, highlighting the characteristic chirp behavior of the signal [5].



(a) Observation plans for the interferometer network (credit to [6]).



(b) Operational and planned ground-based interferometric gravitational-wave detectors (credit to [6]).

Figure 1.2 (Top) Observing plans and operational/planned instrument : The first four observing runs have been completed, yielding a total of 218 recorded events up to the first half of O4 (O4a). Looking ahead, O5 is expected to begin in mid-2027 and continue until approximately mid-2030. (Bottom) IFOs network ground based detectors operational/planned: The figure summarizes the current global network of laser-interferometer observatories, including existing facilities as well as proposed upgrades and next-generation instruments that will extend the astrophysical reach and sensitivity of gravitational-wave astronomy.

1.2 Linearization of Einstein Equations

The Einstein Equations, as every other fundamental equation of physics, are not derived, in the sense that does not exist any epistemological approach that can be applied to set up physical relations among the considered quantities. On the other side, if gravitational action is considered and the variation with respect to the metrics $g_{\mu\nu}$ is computed, then one can obtain the Einstein Equations as

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4}T_{\mu\nu} \quad (1.1)$$

Or, defining the Einstein Tensor as $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu}$. Then¹

¹Do not confuse the Einstein Tensor $G_{\mu\nu}$ with the gravitational constant $G = 6.67 \times 10^{-11} \frac{\text{Nm}^2}{\text{Kg}^2}$

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \quad (1.2)$$

Here $R_{\mu\nu}$ is the Ricci tensor, R is the Ricci scalar (the trace of $R_{\mu\nu}$), $T_{\mu\nu}$ is the stress–energy tensor encoding matter and energy, and $g_{\mu\nu}$ is the spacetime metric. The equations, coupled by the gravitational constant G , elegantly relate geometry and matter: the left-hand side contains purely geometric quantities ($R_{\mu\nu}$, R , and $g_{\mu\nu}$), while the right-hand side contains the physical sources ($T_{\mu\nu}$) that curve spacetime. Despite their compact form, these equations are highly non-linear in the metric, making exact solutions challenging in general.

The bound between space–time’s geometry and mass is described in the Einstein Equations, and thanks to them a lot of physical phenomena can be understood, foreseen and predicted with outstanding precision. The first success of the Einstein’s theory was the explanation of the Mercury’s perihelion precession, unexplained event until the GR’s advent. Since then, this theory continued to show its solidity, describing astrophysical phenomena such as the stellar collapse that marks the birth of the black holes, or even the cosmological description of the expanding universe.

In the present thesis, we will focus on a specific phenomenon that is related to solution of the Einstein Equation on what is called the linearized approximation: the Gravitational Waves (GW). The GWs are astrophysical events that take place when a system presents a variation in time of quadrupole momentum. Physically, they are traveling perturbations, or ripples, of the space–time. Let’s dig further in this direction.

We start recalling how to transform the metric in GR via coordinates transformation. Consider a set of coordinates x^μ , if we want to physically describe what is happening from another coordinate system x'^μ then we have to know how the variables and the metric $g_{\mu\nu}$ change accordingly. This can be formalized with

$$x^\mu \rightarrow x'^\mu(x^\alpha) \quad (1.3)$$

$$g_{\mu\nu} \rightarrow g'_{\mu\nu} = \frac{\partial x^\rho}{\partial x'^\mu} \frac{\partial x^\delta}{\partial x'^\nu} g_{\rho\delta} \quad (1.4)$$

The linearized GR consists in a perturbative solution of the Einstein Equations. If we consider a perturbative term $h_{\mu\nu}$ and a background term $\bar{g}_{\mu\nu}$, then the metric $g_{\mu\nu}$ can be written as

$$g_{\mu\nu} = \bar{g}_{\mu\nu} + h_{\mu\nu}, \text{ with } |h_{\mu\nu}| \ll 1 \quad (1.5)$$

So inserting this expression in the Einstein Equations, these result in

$$G_{\mu\nu}[\bar{g}_{\mu\nu} + h_{\mu\nu}] = \frac{8\pi G}{c^4} T_{\mu\nu}[\bar{g}_{\mu\nu} + h_{\mu\nu}] \quad (1.6)$$

That corresponds to an expansion in h at leading order for $G_{\mu\nu}$ and $T_{\mu\nu}$ as

$$G_{\mu\nu}^{(0)} + G_{\mu\nu}^{(1)} = \frac{8\pi G}{c^4} (T_{\mu\nu}^{(0)} + T_{\mu\nu}^{(1)}) \quad (1.7)$$

We consider as the background term $\bar{g}_{\mu\nu}$ the Minkowski's space-time $\eta_{\mu\nu}$, that consists in the flat metrics. This implies that $G_{\mu\nu}^{(0)}(\eta) = 0$, since the curvature of a flat space-time is zero. By pairing the same order terms in Eq. 1.7 then

$$G_{\alpha\beta}^{(0)} = T_{\alpha\beta}^{(0)} = 0 \quad (1.8)$$

$$G_{\alpha\beta}^{(1)} = \frac{8\pi G}{c^4} T_{\alpha\beta}^{(1)} \quad (1.9)$$

If we now consider the first order term in what is called the trace reversed condition for h , identified by \bar{h} , we have the following set of equations

$$\bar{h}_{\alpha\beta} = h_{\alpha\beta} - \frac{1}{2} \eta_{\alpha\beta} h, \quad (1.10)$$

$$h = Tr(h_{\alpha\beta}) = h^{\alpha\beta} h_{\alpha\beta}, \quad (1.11)$$

$$\eta^{\alpha\beta} \bar{h}_{\alpha\beta} = h - \frac{1}{2} \cdot 4 \cdot h = -h = \bar{h} \quad (1.12)$$

To obtain a simplified expression for the first-order term of Einstein Equations we must impose the Lorentz gauge

$$\partial_\rho \bar{h}_\mu^\rho = 0 \quad (1.13)$$

Then applying Eq. 1.13, we can write from Eq. 1.7 that

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}^{(1)} \quad (1.14)$$

The Lorentz gauge transformation is always possible, since it is always possible to find a set of coordinates that satisfy this condition by considering $x'^\mu = x^\mu + \xi^\mu$. Then we can ask

1. How does $h'_\nu{}^\mu$ transform under Eq. 1.13 ?
2. How does $\bar{h}_{\mu\nu}$ changes via Eq. 1.13 ?
3. Is it possible to choose ξ^μ such that in our coordinates system x'^μ , $\partial_\rho \bar{h}'^\rho{}_\mu = 0$?

The first point is a direct consequence of the metrics transformation rule in Eq. 1.4. In fact by imposing $g'_{\mu\nu} = (\eta'_{\mu\nu} + h'_{\mu\nu})$ and expanding at first order in ξ we find out that $h'_{\mu\nu} = h_{\mu\nu} - 2\partial_\mu \xi_\nu$. The second point follows directly from h' since

$$\bar{h}'_{\mu\nu} = h'_{\mu\nu} - \frac{1}{2} \eta_{\mu\nu} h' = \bar{h}_{\mu\nu} - (\partial_\mu \xi_\nu + \partial_\nu \xi_\mu) + \eta_{\mu\nu} \partial_\alpha \xi^\alpha \quad (1.15)$$

Finally, the third and last point can be obtained applying the differential operator to $\bar{h}'_{\mu\nu}$, then

$$\partial^\mu \bar{h}'_{\mu\nu} = \partial^\mu \bar{h}_{\mu\nu} - \square \xi_\nu - \partial^\mu \partial_\nu \xi_\mu + \partial^\nu \partial_\alpha \xi^\alpha \quad (1.16)$$

And imposing

$$\square \xi_\nu = \partial^\mu \bar{h}_{\mu\nu} \quad (1.17)$$

then $\partial^\mu \bar{h}'_{\mu\nu} = 0$, implying that the Lorentz gauge Eq. 1.13 is a consequence of $\square \xi_\nu = \partial^\mu \bar{h}_{\mu\nu}$. Explicitly, we can always find a solution for ξ_μ in inhomogeneous $\tilde{\xi}$ and homogeneous ξ_ν^{hom} terms:

$$\xi_\nu = \tilde{\xi}_\nu + \xi_\nu^{hom} \quad (1.18)$$

If this solution is applied to Eq. 1.17, then the Lorentz gauge is satisfied. In addition, we can fix the degrees of freedom corresponding to $\square \xi_\nu^{hom} = 0$, that results in 4 conditions. These with the other 4 conditions from the Lorenz gauge fixes 8 constraints. At this point we impose the Transverse Traceless (TT) conditions

- Transverse condition (to the propagation direction i): $\partial^i \bar{h}_{ij} = 0$
- Traceless condition: $\bar{h}^\mu_\mu = 0$

These conditions, TT and the Lorentz gauge Eq. 1.13, are equivalent to

$$h_{0\mu} = 0 \quad (1.19)$$

$$h^i_i = 0 \quad (1.20)$$

$$\partial^i \bar{h}_{ij} = 0 \quad (1.21)$$

That correspond to the following matrix, with the sum of highlighted elements equal to zero.

$$h_{ij} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{22} & h_{23} & h_{24} \\ 0 & h_{32} & h_{33} & h_{34} \\ 0 & h_{42} & h_{43} & h_{44} \end{pmatrix} \quad (1.22)$$

This shows that only the spatial components are relevant to us, and so the Linearized Einstein Equations are

$$\square \bar{h}_{ij} = -\frac{16\pi G}{c^4} T_{ij}^{(1)} \quad (1.23)$$

That is formally the wave equation for the tensor \bar{h}_{ij} in the TT gauge. The solution of these equations will give us the explicit expression of the gravitational waveform signal. It is important to stress that the traceless conditions in this solution need to be imposed, leaving only the physical

degrees of freedom and gauging. That means that we solve these equations for general i, j components and then we impose $h_i^i = 0$ and $\partial^i \bar{h}_{ij} = 0$ via TT conditions. Suppose now that the wave is propagating in some direction \hat{n} (and $|\hat{n}| = 1$), to find the TT components, you need to define a projector operator P_{ij}

$$P_{ij}(\hat{n}) = \delta_{ij} - n_i n_j = P_{ij} \quad (1.24)$$

This operator ensures the transverse component projection of a vector and in doing so it isolates the two independent polarization of the gravitational wave. Now we can construct a projector onto the TT components of a generic tensor defining Λ_{ijkl}

$$\Lambda_{ijkl}(\hat{n}) = P_{ik}(\hat{n})P_{jl}(\hat{n}) - \frac{1}{2}P_{ij}P_{kl} \quad (1.25)$$

This Λ_{ijkl} is symmetric, transverse and traceless with respect to each index pair. It can be applied to any symmetric tensor \bar{h}_{ij} to extract the physical-relevant transverse-traceless parts

$$h_{ij}^{TT} = \Lambda_{ijkl} \bar{h}^{kl} \quad (1.26)$$

Let's now solve the Linearized Einstein Equations 1.14, or equivalently the Eq. 1.23, in vacuum by imposing $T_{ij}^{(1)} = 0$. We are looking for a solution of this equation

$$\square \bar{h}_{ij} = 0 \quad (1.27)$$

that is the homogeneous wave-equation, so the solution will be wave-like and in the frequency space κ it will be

$$\bar{h}_{ij}(x) = \int d\kappa^4 (\bar{h}_{ij}(\kappa) e^{i\kappa_\mu x^\mu}) \quad (1.28)$$

The solution is the plane-wave equation

$$\bar{h}_{ij}(x) = \epsilon_{ij} e^{ik \cdot x} \quad (1.29)$$

Plugging this ansatz into Eq. 1.27 results then in the algebraic equation $\kappa^2 \bar{h}_{ij} = 0$, implying $\kappa^\mu \kappa_\mu = \kappa^2 = 0$ for non-trivial solution — i.e. gravitational waves travel at speed of light since it is a null quadri-vector. Since $\kappa = (\omega, \vec{\kappa})$ and by $\kappa^2 = 0$, it follows $|\vec{\kappa}|^2 = \omega^2$. For propagation along \hat{z} axis then $\vec{\kappa} = (0, 0, \pm\omega)$. It follows that the Lorentz gauge conditions $\partial^j \bar{h}_{ij} = 0$ translates in $\kappa^j \bar{h}_{ij}(\kappa) = 0$, this means that

$$\omega \bar{h}_{3j} = 0 \rightarrow \bar{h}_{3j} = 0 \quad (1.30)$$

So the \bar{h}_{ij} has the following structure

$$\bar{h}_{ij}^{TT} = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (1.31)$$

It is evident that the components are solely in the \hat{x} and \hat{y} directions reflecting the transversality of the gravitational wave propagating along the \hat{z} direction. The $-h_+$ term has been imposed by hand due to traceless condition. At this point, since h_{ij}^{TT} is known, we can study how the perturbations enters in the metrics ds . In doing so, we remark that the h solution has been found and we are working within the TT components via projection procedure described by Eq. 1.26, and that h_+ and h_\times represent plane-wave solution of $h_{+,\times}(t, \pm z) = A_{+,\times} e^{i\omega t \pm z}$. The metrics ds results in

$$ds^2 = g_{\mu\nu} dx^\mu dx^\nu = (\eta_{\mu\nu} + h_{\mu\nu}) dx^\mu dx^\nu = \\ + (1 + h_+(t \pm z)) dx^2 + (1 - h_+(t \pm z)) dy^2 + 2h_\times(t \pm z) dx dy \quad (1.32)$$

With h_+ and h_\times depending on the source of the GW. Now we can ask ourselves, how does the GW affects the particles? Consider now a free test mass at rest (free falling particle), we have to consider the geodesic equation

$$\frac{d^2 x^\mu}{d\tau^2} + \Gamma_{\alpha\beta}^\mu \frac{dx^\alpha}{d\tau} \frac{dx^\beta}{d\tau} = 0. \quad (1.33)$$

Focusing only on the spatial components $\mu = i$, we then ask: does a particle initially at rest will show an acceleration due to the passage of gravitational wave? Imposing initial conditions $\frac{dx^0}{d\tau} = 1$ and $\frac{dx^i}{d\tau} = 0$, the spatial equation for acceleration is

$$\frac{dx^i}{d\tau} = -\Gamma_{00}^i \quad (1.34)$$

The spatial components will give us 0, in fact in TT gauge the Christoffel's symbols are

$$\Gamma_{00}^i = \frac{1}{2} \eta^{ij} (2 \underbrace{\partial_0 h_{j0}^{TT}}_0 - \underbrace{\partial_i h_{00}^{TT}}_0) = 0 \quad (1.35)$$

so the particle does not move! The particle stays at rest and does not change coordinate location. This because the TT gauge represents a coordinate system that is comoving with freely-falling particle. We need to look at one invariant quantity. Let's define the invariant distance between two events. In 1-D let's pick $(t, x_1, 0, 0)$ and $(t, x_2, 0, 0)$. Then we define $L = x_1 - x_2$. The metrics is affected at the passage of the GW in the following way

$$\Delta s_x^2 = g_{\mu\nu} \Delta x^\mu \Delta x^\nu = g_{11} (x_1 - x_2)^2 = \\ = (\eta_{11} + h_{11}) L^2 = (1 + h_+ \cos(\omega t)) L^2 \quad (1.36)$$

Then

$$\Delta s_x = L[1 + h_+ \cos(\omega t)]^{\frac{1}{2}} \sim L(1 + \frac{1}{2}h_+ \cos(\omega t)) \quad (1.37)$$

So an oscillation is measurable in the distance! Let's now consider a classical example, the ring of test masses (particles) in $\hat{x} - \hat{y}$ plane. The displacement along the i -th direction is

$$\delta s^i(t) = \delta s^i(o) + \frac{1}{2}h_{ij}^{TT} \delta s_j(o) \quad (1.38)$$

where $\delta s_j(o)$ is invariant (invariant distance along j direction). So, suppose that only the h_+ polarization affects the particles, then the tensor becomes

$$\bar{h}_{ij}^{TT} = \begin{pmatrix} h_+ & 0 & 0 \\ 0 & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix} \cdot \sin(\omega t \pm z) \quad (1.39)$$

the displacement along the directions are

$$\Delta z \rightarrow \delta s^3(0) = 0 \text{ due to } h_{i3}^{TT} = 0 \quad (1.40)$$

$$\Delta x \rightarrow \delta s^1(t) = \delta s^1(0) + \frac{1}{2}(h_+ \sin(\omega t)) \Big|_{z=0} \delta s^1(0) = \delta x(0) [1 + \frac{1}{2}(h_+ \sin(\omega t))] \quad (1.41)$$

$$\Delta y \rightarrow \delta s^2(t) = \delta y(0) [1 - \frac{1}{2}(h_+ \sin(\omega t))] \quad (1.42)$$

Combining \hat{x} and \hat{y} together, if the particle at $t = 0$ is on circle of radius $R = 1$, $\delta x(0)^2 + \delta y(0)^2 = 1$ ad time t instead it describes an ellipse (see Fig. 1.3 for reference).

$$[1 + \frac{1}{2}h_+ \sin(\omega t)]^2 + [1 - \frac{1}{2}h_+ \sin(\omega t)]^2 = 1 \quad (1.43)$$

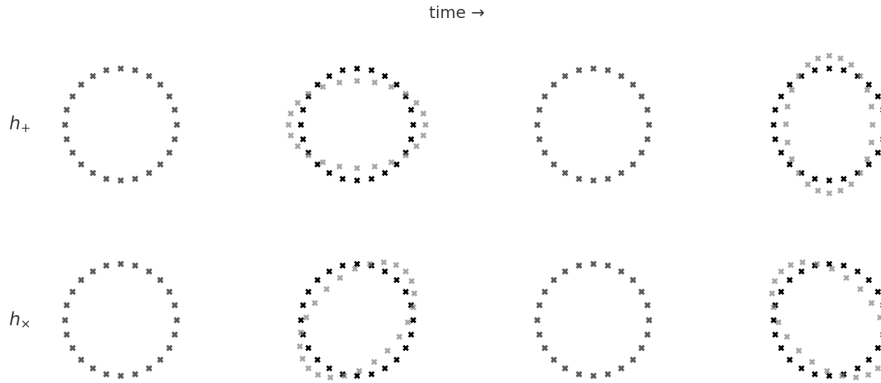


Figure 1.3 Motion of test particles in + and x polarization polarization

1.3 The generation mechanism in linearized theory

Now we will no longer consider the vacuum hypothesis and study the case where matter enters in the equations. We then aim to solve the inhomogeneous Linearized Einstein Equations

$$\square \bar{h}_{ij} = -\frac{16\pi G}{c^4} T_{ij}^{(1)} \quad (1.44)$$

The source is spatially localized with a characteristic-size L , and with no spherical symmetry assumption at this stage, as represented in Fig 1.4. The solution of Eq. 1.44 can be found by the retarded Green Function in flat spacetime

$$\bar{h}_{ij}(x, t) = \frac{4\pi}{c^4} \int_{source} d^3y \frac{T_{ij}(t - \frac{|\vec{x}-\vec{y}|}{c}, \vec{y})}{|\vec{x}-\vec{y}|} \quad (1.45)$$

Now we assume also that the observer is located far away from the source $R \gg L$. This allows us to apply a Taylor expansion of the solution in $\frac{L}{R}$. While this is straightforward for the denominator, the numerator requires a further passage. The expansion in $\frac{L}{R}$ is equivalent to $\frac{v}{c}$. In fact

$$\begin{aligned} T_{ij}(t - \frac{|\vec{x}-\vec{y}|}{c}, \vec{y}) &= T_{ij}^{(1)}(t - \frac{R}{c}, \vec{y}) + \\ &+ \frac{\vec{y} \cdot \hat{n}}{c} \partial_t T_{ij}^{(1)}(t - \frac{R}{c}, \vec{y}) \end{aligned} \quad (1.46)$$

Let's now show that $\frac{\vec{y} \cdot \hat{n}}{c}$ is effective of order $\frac{v}{c} \ll 1$. Given t_c the characteristic timescale of the source's internal dynamic, then

$$\partial_t T_{ij}^{(1)} \sim \frac{T_{ij}}{t_c} \quad (1.47)$$

$$\frac{\vec{y} \cdot \hat{n}}{ct_c} \sim \frac{L}{ct_c} \sim \frac{L}{t_c} \frac{1}{c} \sim \frac{v}{c} \ll 1 \quad (1.48)$$




Figure 1.4 Geometry of \vec{x} , \vec{y} e $\vec{x}-\vec{y}$ inside the body of typical length L .

The first order correction is suppressed by a factor $o(\frac{v}{c})$, validating the non-relativistic approximation. It is possible then to describe the dynamic of the source rely solely on Newtonian physics. Under these assumptions the solution to the linearized Einstein Equations becomes

$$\bar{h}_{ij} = \frac{4G}{c^4 R} \int d^3y T_{ij}^{(1)}(t - \frac{R}{c}, \vec{y}) + O(\frac{L^2}{R}) \quad (1.49)$$

It is now convenient to introduce the stress-moments of the energy-momentum tensor S_{ij} defined by

$$S_{ij} = \int d^3y T_{ij}^{(1)}(\bar{t}, \vec{y}) \quad (1.50)$$

$$S_{ij,k} = \int d^3y T_{ij}^{(1)}(\bar{t}, \vec{y}) y_k \quad (1.51)$$

⋮

where $\bar{t} = t - \frac{R}{c}$. These moments will play a central role in constructing the multipole expansion of the gravitational wave field and, in particular, in deriving the quadrupole formula. Then equation Eq. 1.49 at first order can be expressed as

$$\bar{h}_{ij} = \frac{4G}{c^4 R} S_{ij} \quad (1.52)$$

This expression becomes particularly useful if we define the linear momentum P and the energy momentum density I .

$$I = \frac{1}{c^2} \int d^3y T_{00}^{(1)}(\bar{t}, \vec{y}) \quad (1.53) \quad P^i = \frac{1}{c} \int d^3y T_{0i}^{(1)}(\bar{t}, \vec{y}) \quad (1.55)$$

$$I^i = \frac{1}{c^2} \int d^3y T_{00}^{(1)}(\bar{t}, \vec{y}) x^i \quad (1.54) \quad P^{i,j} = \frac{1}{c} \int d^3y T_{0i}^{(1)}(\bar{t}, \vec{y}) x^j \quad (1.56)$$

⋮

⋮

The continuity equation, that states the local conservation of energy-momentum expressed by $\partial^\mu T_{\mu\nu}^{(1)} = 0$, allows us to relate the stress moment S^{ij} to the energy density quadrupole moment I . Integrating by part the continuity equation one obtains

$$S^{ij} = \frac{1}{2} \ddot{I}^{ij} \quad (1.57)$$

Inserting this expression in \bar{h}_{ij} we can express the quadrupole formula for gravitational field in linearized gravity, obtaining the well known result that the GW production is related to the second-order time derivative of the quadrupole moment I :

$$\bar{h}_{ij} = \frac{2G}{c^4 R} \partial_t^2 \underbrace{\int d^3y T_{00}(\bar{t}, \bar{y}) y_i y_j}_{\text{quadrupole moment } I_{ij}} = \frac{2G}{c^4 R} \ddot{I}_{ij} \quad (1.58)$$

Up to now, we have treated the gravitational wave as a perturbative term h on the flat Minkowski space-time $\eta_{\mu\nu}$. Now we broaden our setting and consider a generic curved background metric $\bar{g}_{\mu\nu}$. We will see that this consideration is fundamental for properly accounting the energy and momentum carried by GWs, that is a key conceptual point of the theory. We assume that the characteristic curvature scale of the background is L_b , while the GW's length is λ_{GW} , with hierarchy $\lambda_{\text{GW}} \ll L_b$. In frequency space the relation translates to $f_b \ll f_{\text{GW}}$. This scales' separation justifies the second order expansion of Einstein Equations in h , since only at second order the wave carries energy. Recall the Einstein Field Equations

$$R_{\mu\nu} = 8\pi G(T_{\mu\nu} - \frac{1}{2}g_{\mu\nu}T) \quad (1.59)$$

We now expand the Ricci tensor in left side (LHS) in a series of h , clearly the superscripts identify the order of expansion.

$$R_{\mu\nu} = \bar{R}_{\mu\nu}^{(0)} + \bar{R}_{\mu\nu}^{(1)} + \bar{R}_{\mu\nu}^{(2)} \quad (1.60)$$

the lowest term $\bar{R}_{\mu\nu}^{(0)}$ corresponds to the background curvature. Slow-varying metrics is associated to only low frequencies modes. The $\bar{R}_{\mu\nu}^{(1)}$ is linear in h and captures the GW perturbation term, so only high frequencies contents are present. Finally the $\bar{R}_{\mu\nu}^{(2)}$ is quadratic in h , then it contains both high and low frequencies components.

We now decompose the Einstein Equations into low and high frequency parts, exploiting this scale separation.

$$\begin{cases} \text{LF: } \bar{R}_{\mu\nu}^{(0)} = -\bar{R}_{\mu\nu}^{(2)}|_{\text{low}} + (T_{\mu\nu} - \frac{1}{2}g_{\mu\nu}T)|_{\text{low}} \\ \text{HF: } \bar{R}_{\mu\nu}^{(1)} = -\bar{R}_{\mu\nu}^{(2)}|_{\text{high}} + (T_{\mu\nu} - \frac{1}{2}g_{\mu\nu}T)|_{\text{high}} \end{cases} \quad (1.61)$$

the high-frequency (HF) equation in the decomposition above allows us to define an effective stress energy tensor for gravitational waves by isolating terms quadratic in perturbation h . This tensor encodes the energy, momentum, and stress carried by gravitational waves, even though gravity is inherently non-linear, as Einstein Equations describe. Formally, we define the effective gravitational wave stress-energy tensor (often associated with the Landau–Lifshitz pseudotensor) as:

$$t_{\mu\nu}^{\text{GW}} = -\frac{1}{8\pi G} (R_{\mu\nu}^{(2)} - \frac{1}{2}\bar{g}_{\mu\nu}R^{(2)}) \quad (1.62)$$

To extract physically meaningful behavior from the low-frequency components we average over several periods $T = \frac{1}{f_{\text{GW}}}$, or equivalently among several wavelengths. Denoting this average as $\langle \cdot \rangle$ we find that in the TT gauge

$$t_{\mu\nu}^{\text{GW}} = \frac{c^4}{32\pi\text{G}} \langle \partial_\mu h_{\alpha\beta} \partial_\nu h^{\alpha\beta} \rangle \quad (1.63)$$

$$t_{00}^{\text{GW}} = \frac{c^4}{32\pi\text{G}} \langle \partial_0 h_{\alpha\beta}^{TT} \partial_0 h^{\alpha\beta,TT} \rangle = \frac{c^4}{16\pi\text{G}} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle \quad (1.64)$$

This result shows clearly that GWs do carry energy, linear momentum and angular momentum. In particular

- h_{00} : related to energy density (gauge invariant at linear order),
- h_{0i} : associated with linear momentum density,
- h_{ij} : associated with stress and angular momentum flux.

The effective tensor $t_{\mu\nu}^{\text{GW}}$ allows us to compute observable effects such as energy loss due to gravitational wave emission, most notably in binary systems (e.g., the inspiral of neutron stars or black holes). We can compute the energy flux carried by a GW, that is the energy transmitted per unit time through a unit surface at large distance from the source. Consider a spatial volume V bounded by a surface S . The conservation of the effective energy-momentum tensor, in this case the Landau-Lifshitz pseudotensor, yields

$$\int_V d^3x (\partial_0 t^{00} + \partial_i t^{i0}) = 0 \quad (1.65)$$

The energy contained in the volume is:

$$E_V = \int_V d^3x \partial_0 t^{00} = 0 \quad (1.66)$$

Taking the time derivative and applying the divergence theorem, we obtain:

$$\frac{1}{c} \frac{dE_V}{dt} = - \int_V d^3x \partial_i t^{i0} = - \int_S dA n_i t^{0i} \quad (1.67)$$

Where n^i is the outward-pointing unit normal vector to the surface element dA . The decrease of E_V , indicates that energy is radiated outward by the passing gravitational wave. Therefore, the energy flux (power per unit area) carried by the wave is:

$$\frac{dE_V}{dt dA} = \frac{c^3}{16\pi\text{G}} \langle \dot{h}_+^2 + \dot{h}_\times^2 \rangle \quad (1.68)$$

1.4 The inspiral of compact objects

One of the most significant source of GWs in astrophysics is the inspiral and merger of compact binary systems such as binary black holes (BBHs) or binary neutron stars (BNSs). These systems

emit gravitational radiation during their orbital evolution, and their signal can be accurately modeled by solving the Einstein Equations in the appropriate regimes. In this section we consider the inspiral phase, during which two objects orbit each other in nearly circular, bounded Keplerian motion, and gravitational radiation causes a slow decay of the orbital radius. For simplicity we will ignore spin effects and consider purely Newtonian circular orbit. We work in the center of mass frame and circular orbits of radius a . The relevant quantities are

- Center of mass: $m_1\vec{y}_1 + m_2\vec{y}_2 = 0$; $\vec{y}_1 = \frac{m_2}{M}\vec{r}$; $\vec{y}_2 = -\frac{m_1}{M}\vec{r}$; $M = m_1 + m_2$
- Relative position: $\vec{r} = \vec{y}_1 - \vec{y}_2$, $\vec{r} = (a\cos(\Omega t), a\sin(\Omega t))$; $\Omega = \sqrt{\frac{GM}{a^3}}$, where Ω is the orbital angular frequency from Kepler's third law.

The leading order GWs emission arises from the second time derivative of the mass quadrupole moment. For point particles, the quadrupole moment is given by

$$\begin{aligned} I_{ij} &= \int d^3y \overbrace{T_{00}}^{\text{masses}} \left(t - \frac{R}{c}, \vec{y}\right) y_i y_j \\ &= m_1 y_i^1 y_j^1 + m_2 y_i^2 y_j^2 = \mu r_i r_j \end{aligned} \quad (1.69)$$

Where $\mu = \frac{m_1 m_2}{M}$ is the reduced mass and r_i, r_j are the components of the relative separation vector. This form highlights that the time dependence of the gravitational wave is fully encoded in the orbital motion of the binary through $\vec{r}(t)$. Let's analyze the explicit form of the quadrupole moment. Given the relative position in the orbital plane

$$\vec{r}(t) = a(\cos \Omega t, \sin \Omega t, 0) \quad (1.70)$$

We compute

$$I_{xx}(t) = \mu a^2 \cos^2(\Omega t), \quad (1.71)$$

$$I_{yy}(t) = \mu a^2 \sin^2(\Omega t), \quad (1.72)$$

$$I_{xy}(t) = \mu a^2 \cos(\Omega t) \sin(\Omega t). \quad (1.73)$$

Taking the second time derivatives, we obtain:

$$\ddot{I}_{xx}(t) = -2\mu a^2 \Omega^2 \cos(2\Omega t) = -2\mu \frac{GM}{a} \cos(2\Omega t), \quad (1.74)$$

$$\ddot{I}_{yy}(t) = +2\mu \frac{GM}{a} \cos(2\Omega t), \quad (1.75)$$

$$\ddot{I}_{xy}(t) = -2\mu \frac{GM}{a} \sin(2\Omega t), \quad (1.76)$$

Where we used $\Omega^2 = \frac{GM}{a^3}$. Inserting this in the linearized Einstein solution for the metric perturbation h in TT gauge at leading order we find:

$$h_{ij}^{TT}(t, \vec{x}) = \frac{2G}{c^4 R} \ddot{I}_{ij}(t - R/c). \quad (1.77)$$

Thus, the final expression for the TT metric perturbation is:

$$h_{ij}^{TT}(t, x) = \frac{2G}{c^4 R} 2\mu \frac{GM}{a} \begin{pmatrix} \cos(2\Omega\bar{t}) & \sin(2\Omega\bar{t}) & 0 \\ \sin(2\Omega\bar{t}) & -\cos(2\Omega\bar{t}) & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad (1.78)$$

This matrix shows that the gravitational wave has h_+ and h_\times components oscillating at frequency 2Ω . We stress that this double frequency incorporate the quadrupole nature of the source.

Now, since the gravitational wave frequency is twice the orbital frequency, $f_{\text{GW}} = 2\Omega$, then we can express the orbital separation a in terms of f_{GW} using the third Kepler's law.

$$f_{\text{GW}} = \frac{1}{\pi} \sqrt{\frac{GM}{a^3}} \quad (1.79)$$

It is convenient to introduce now a key physical quantity: the chirp mass M_c defined as

$$M_c = \mu^{\frac{3}{5}} M^{\frac{2}{5}} = \frac{(m_1 m_2)^{\frac{3}{5}}}{(m_1 + m_2)^{\frac{1}{5}}} \quad (1.80)$$

This mass parameter governs both amplitude and frequency evolution of the emitted gravitational wave. The amplitude of the tensor h in terms of f_{GW} and M_c is then

$$h_0 = \frac{2G}{c^4 R} 2\mu \frac{GM}{a} = \frac{\sqrt{2}}{c^4 R} (2\pi f_{\text{GW}})^{\frac{2}{3}} (GM_c)^{\frac{5}{3}} \quad (1.81)$$

This expression states that the amplitude grows as the frequency increases. This is the physical origin of the 'chirp' behavior in gravitational wave signals produced by inspiraling objects. We can now calculate the energy radiated in circular orbits. Let's consider a system with circular radius a , and exploiting the known relation between f_{GW} and Ω , the two polarizations in the TT gauge (we suppress the superscript for simplicity) are then:

$$h_+(t) = \frac{\sqrt{2}}{c^4 R} (GM_c)^{5/3} (2\pi f_{\text{GW}})^{2/3} \cos(2\pi f_{\text{GW}} \bar{t}), \quad (1.82)$$

$$h_\times(t) = \frac{\sqrt{2}}{c^4 R} (GM_c)^{5/3} (2\pi f_{\text{GW}})^{2/3} \sin(2\pi f_{\text{GW}} \bar{t}), \quad \bar{t} = t - \frac{R}{c} \quad (1.83)$$

These are the leading-order quadrupole expressions for the gravitational waveforms from a non-spinning, circular, compact binary inspiral, and form the foundation for both waveform modeling and

matched filtering in GW data analysis. To understand the temporal evolution of the system, we take into account the energy balance equation.

$$\frac{dE_{orbit}}{dt} = -P_{GW} \quad (1.84)$$

This equation relates the orbital energy rate loss due to gravitational waves power emission. From the Newtonian mechanics the orbital energy for binary system separate by distance a is $E_{orbit} = -\frac{GM}{2a}$. Differentiating with respect to time and recalling $\mu M = M_c^{5/3} M^{1/3}$, we find

$$\frac{dE}{dt} = \frac{GM_c}{2a} \dot{a} \propto \frac{\dot{f}_{GW}}{f_{GW}^{1/3}} \quad (1.85)$$

On the other hand, the energy emitted by gravitational wave radiation, computed from the Landau-Lifshitz pseudotensor and the quadrupole formula, is

$$P_{GW} = \frac{32}{5} \frac{c^5}{G} \left(\frac{GM_c}{4c^3} f_{GW} \right)^{10/3} \quad (1.86)$$

Equating then the terms in Eq. 1.84 we finally obtain the chirp formula

$$\dot{f}_{GW} = \frac{96}{5} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} f_{GW}^{11/3} \quad (1.87)$$

This first-order differential equation governs the frequency evolution of the system. By integrating this expression we find

$$f(t)_{GW} = \left[\frac{8}{3} \pi^{8/3} \left(\frac{GM_c}{c^3} \right)^{5/3} (t_c - t) \right]^{-3/8} \quad (1.88)$$

Which diverges for $t = t_c$. The time t_c is interpreted as coalescence time, that is the moment when the two objects merge and the current assumptions under which the solutions are obtained are no longer valid. The rapid increase of frequency and amplitude as $t \rightarrow t_c$ justifies the 'chirp' behavior of gravitational wave signal (see Fig. 1.5).

The maximum frequency introduces a lower bound for orbital separation a . For instance, considering $f_{max} \sim 1\text{kHz}$, and a binary system $m_1 = m_2 = 1.4M_\odot$, one finds $a_{min} \sim 33\text{km}$, it is evident that such compact objects can approach extremely close distances before merging. All the current computation involved uniquely the variation of the frequency f_{GW} , the next step is to capture the full time-dependence of the waveform, which imply modeling the time-varying phase $\phi(\bar{t})$, derived by the orbital frequency $\Omega(t)$. The waveform polarization then becomes

$$h_+(t) = \frac{\sqrt{2}}{c^4 R} f_{\text{GW}}^{\frac{2}{3}} (GM_c)^{\frac{5}{3}} \cos(\phi(\bar{t})) \quad (1.89)$$

$$h_\times(t) = \frac{\sqrt{2}}{c^4 R} f_{\text{GW}}^{\frac{2}{3}} (GM_c)^{\frac{5}{3}} \sin(\phi(\bar{t})), \quad \bar{t} = t - \frac{R}{c} \quad (1.90)$$


All calculations presented so far rely on weak-field approximation and Minkowski flat background. However, as the inspiral binary approaches the final stage, gravitational field becomes strong and this approximation breaks down, and numerical relativity and Black-hole perturbation theory are required [34, 35]. One way to estimate the transition to this non-linear regime is through the Inner Most Stable Circular Orbit (ISCO), beyond which stable orbits are no more allowed. For non spinning systems this is defined as

$$a_{\text{ISCO}} = \frac{6GM}{c^2} \quad (1.91)$$

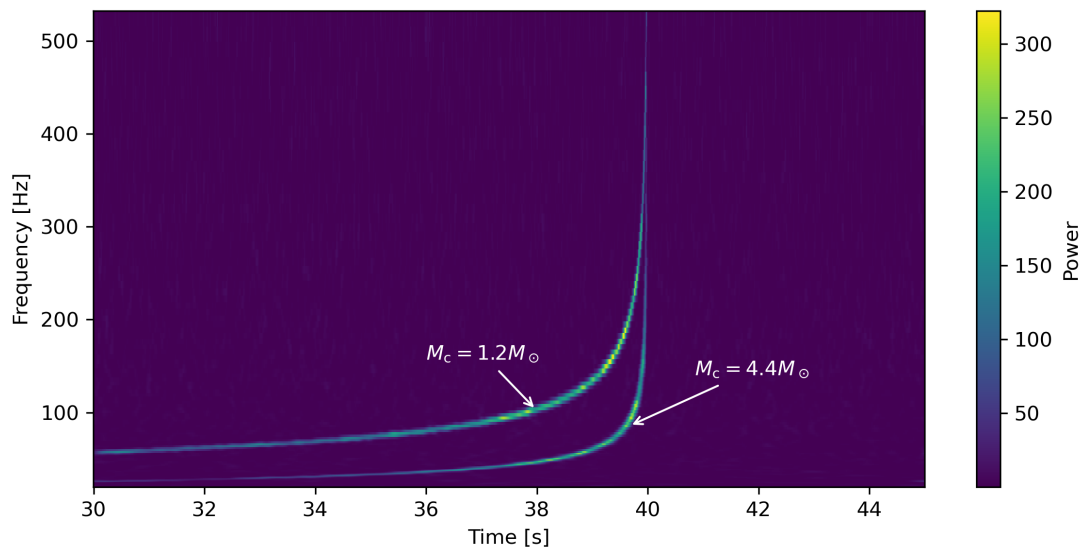
If we then assume that the merge happens near this separation, we can estimate the merger frequency as

$$f_{\text{merger}} = 2\sqrt{\frac{GM_c}{a_{\text{ISCO}}^3}} \sim \frac{1}{GM_c} \quad (1.92)$$

This provides a rough but valuable upper limit on the detectable gravitational wave frequency from compact binary inspirals before coalescence.



(a) Signals generated from a Binary-Neutron-Star ($M_c = 1.2M_\odot$) injected at 20Mpc (top) and a Binary-Black-Hole ($M_c = 4.4M_\odot$) injected at 100Mpc (bottom).



(b) Time-frequency representation (Q-transform spectrogram) for the same data stretch, showing the binary neutron star and binary black hole signals [5].

Figure 1.5 Different source types, assuming the same starting frequency, produces signals with different durations: hundreds of seconds for a BNS system, and few tens of seconds for a BBH. In this example the merger time is $t_c = 40s$. The spectrogram also highlights the different chirp behavior of the two systems.

1.5 Main observational results

Since the first direct detection of gravitational waves (GWs) from GW150914, the LIGO–Virgo–KAGRA (LVK) Collaboration has reported 218 GW signals from compact binary coalescences up to the first half of the fourth observing run (O4a). Among these, some events have had particularly high scientific impact in a variety of research areas, from testing the robustness of General Relativity to constraining astrophysical processes such as nucleosynthesis —especially as a dominant formation channel of heavy elements [36] — and the properties of matter at supranuclear densities. In particular, the event GW170817 [10] is a cornerstone of what is now called multi-messenger astrophysics. This exceptional event corresponds to the first observation, by the LIGO–Virgo interferometers, of a binary neutron star inspiral, with a total mass $M \in [2.79, 3.29] M_{\odot}$ and component masses in the range $1.17\text{--}1.6 M_{\odot}$. This GW signal was followed by the short gamma-ray burst GRB 170817A, which allowed stringent constraints to be placed on the difference between the speed of light c and the propagation speed of gravitational waves [37]. This extraordinary event also identified a source consistent with kilonova emission from a neutron star merge [38]. Furthermore, this event enabled systematic studies of the neutron star equation of state [39] and has been used as a standard siren for measuring the Hubble constant [40].

Another honorable mention is GW190521, a binary black-hole (BBH) merger with source-frame component masses $m_1 = 85^{+21}_{-14} M_{\odot}$ and $m_2 = 66^{+17}_{-18} M_{\odot}$ and a remnant of $M_f = 142^{+28}_{-16} M_{\odot}$ [11]. The primary mass lies within the so-called (pulsational) pair-instability supernova mass gap ($\sim 65\text{--}120 M_{\odot}$), making GW190521 a key event for understanding the limits of stellar evolution and the formation of black holes in this region of parameter space [11]. The remnant mass places it in the intermediate-mass black-hole (IMBH) regime, providing one of the first gravitational-wave observations of an IMBH. Several formation scenarios have been proposed for such a system, including dynamical assembly via three-body encounters in young massive star clusters [41] and mergers of stellar-mass black holes embedded in an active galactic nucleus (AGN) disc, possibly associated with the flare ZTF19abanrhr in AGN J124942.3+344929 [42]. These models highlight GW190521 as a benchmark event for studying both dense stellar environments and AGN-assisted BBH formation.

Systems that present particularly interesting scientific implications are neutron star–black hole (NSBH) binaries. The first confident detections of such systems, GW200105 and GW200115 [43], involved compact objects with component masses of order $(m_{\text{BH}}, m_{\text{NS}}) \simeq (8.9, 1.9) M_{\odot}$ and $(5.7, 1.5) M_{\odot}$, respectively, with the secondaries consistent with neutron stars at high probability [43]. These events opened a new observational window between the binary neutron-star and binary black-hole populations, providing constraints on the mass and spin distributions of NSBH systems and on the stellar evolution pathways that produce them. Although no secure electromagnetic counterparts were observed for GW200105 and GW200115, the lack of bright emission is itself informative, suggesting that, for the measured mass ratios and black-hole masses, the neutron stars were likely swallowed whole inside the black-hole horizons with little or no ejecta. Building on these detections, population studies have investigated the mass and spin distributions and the expected rates of NSBH mergers, as

well as their multimessenger prospects, by combining the LVK catalogs with population-synthesis and kilonova/GRB models [44, 45].

Finally, we mention the most massive binary black-hole system detected so far, GW231123 [46]. This event is consistent with the merger of two black holes with source-frame masses of order $\sim 130M_{\odot}$ and $\sim 100M_{\odot}$, producing a remnant with a mass exceeding $225M_{\odot}$ [46]. At least one of the components lies in the upper pair-instability mass gap, and both black holes are inferred to have high dimensionless spins, making GW231123 particularly challenging to explain within standard isolated binary evolution scenarios. As a result, it has already motivated detailed studies of alternative formation channels, such as hierarchical mergers and stellar collisions in dense star clusters, as well as more exotic possibilities including primordial black holes [47]. Given its large inferred mass and distance, GW231123 is also an excellent laboratory for investigating gravitational lensing of gravitational waves: current analyses explore whether strong lensing or microlensing could bias the inferred source parameters, although no compelling evidence for lensing has been found so far [46].

1.6 Sources and detection

Those detections relies on the interferometric techniques to measure tiny space-time ripples induced by the gravitational-wave passage. The operational frequency band of current ground-based interferometer spans approximately in $10\text{Hz} \div 5\text{kHz}$ interval, which is dictated by both seismic noise at low frequencies and quantum shot noise at high frequencies. The sensitivity of an interferometer to a GW of wavelength λ_{GW} is characterized by the ratio

$$f_L \sim \frac{L}{\lambda_{\text{GW}}} \quad (1.93)$$

Where L is the interferometer's arm length. For LIGO and Virgo detectors with $L \sim \text{km}$ and $\lambda_{\text{GW}} \sim 10^3\text{km}$ we typically have $f_L \ll 1$, that means that these instruments are sensitive to long-wavelength signal relative to their size. Looking ahead, gravitational wave detection will extend across the full GW spectrum through complementary techniques and observatories. These include:

1. 1. Pulsar Timing Array (PTA) – 10^{-9} – 10^{-7} Hz, $L \sim 10^{17}$ km, $fL \gg 1$ [48, 49].
2. 2. Laser Interferometer Space Antenna (LISA) – 10^{-4} – 1 Hz, $L \sim 10^6$ km, $fL \sim 1$ [50, 51].
3. 3. Einstein Telescope (ET), Cosmic Explorer (CE) – 1 – 10^4 Hz, $L \sim 10$ km, $fL < 1$ [52–54].

Together, these detectors will provide coverage over nearly the entire gravitational wave spectrum, enabling multi-band GW astronomy and deeper insight into the structure and evolution of the universe.

Theoretical understanding of the GW waveforms from different sources is required for both detection and interpretation. This task may be complicated, since the Einstein Equations are inherently non-linear and analytically intractable for general systems. To make progress, a perturbative approach, valid when the space-time curvature is weak and the gravitational field can be expressed as a small

deviation from a known background metric. We saw that GW are associated to perturbation of the space-time ruled by the Einstein Equations

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu,\nu} \quad (1.94)$$

$\underbrace{\hspace{1.5cm}}_{\substack{10^{-43} \\ \text{kg}\cdot\text{m}\cdot\text{s}^{-2}}}$

The coupling constant $\frac{8\pi G}{c^4}$ is extremely small in SI units, implying that only sources with immense stress-energy densities produce detectable space-time perturbations. Consequently, measurable GW can be originated from highly energetic astrophysical sources. For instance, the Earth in its revolution around the Sun emits in form of GW a radiated power $\sim 200 \frac{\text{J}}{\text{y}}$, that is negligible. By contrast, the GW150914 event released an energy approximately of 10^{48}J in 0.2s, with initial masses $m_1 \sim 36M_\odot$ and $m_2 \sim 29M_\odot$ and final $m_{\text{final}} \sim 62M_\odot$, with $3M_\odot$ emitted in gravitational radiation, making it detectable from ground-based interferometers despite the weak coupling in Eq. 1.94. Such enormous release of energy justifies the perturbative expansion in analyzing GW signals: even though the emitted power is vast, the strain remains small at the detector's location, allowing the linearized gravity approach. Then we can ask: for how long our system can be detected? The emission frequency f_{GW} for an inspiral signal can be approximated as Eq. 1.88

$$f_{\text{GW}}(t) \sim \left(\frac{GM_c}{c^3} \right)^{-\frac{5}{3}} \tau^{-\frac{3}{8}} \quad (1.95)$$

with $\tau = t_c - t$ and t_c is the coalescence time. Since for $t \rightarrow t_c$, $f_{\text{GW}} \rightarrow \infty$ then the inspiral phase must be truncated at a finite frequency value by imposing a cut-off frequency f_{merge} . This frequency marks the transition to inspiral phase to merger.

The f_{merge} depends on the total mass, then for different astrophysical binary systems we will have cut-off frequency as summarized in the following table:

<i>f</i> merge for various astrophysical systems		
BNS	BBH	SMBH^a
<ul style="list-style-type: none"> • $m_1 \sim m_2 \sim 1.4M_\odot$ • $f_{\text{merge}} \sim 1000 \text{ Hz}$ 	<ul style="list-style-type: none"> • $m_1 \sim m_2 \sim 35M_\odot$ • $f_{\text{merge}} \sim 60 \text{ Hz}$ 	<ul style="list-style-type: none"> • $m_1 \sim m_2 \sim 10^6 M_\odot$ • $f_{\text{merge}} \sim 10^{-3} \text{ Hz}$
<p>^aSMBH stands for Super Massive Black Holes and consists in singularities with mass above $10^5 M_\odot$, up to billions of M_\odot. Such astrophysical objects can be found at the center of galaxies and are responsible for powering Active Galactic Nuclei and quasars [55]. Also, emission of GW from such objects is predicted [56].</p>		

The measurable duration of a signal depends on the detector's sensitive frequency band, typically bounded at low frequencies by the instrument's sensitivity limit f_{low} ($\sim 10 - 20 \text{ Hz}$) and at high frequencies by the merger frequency f_{merge} . The observable interval therefore spans approximately the range $f_{\text{low}} - f_{\text{merge}}$. The corresponding time before the merger can be estimated as

$$\tau \sim f_{\text{low}}^{-8/3} \left(\frac{GM_c}{c^3} \right)^{-5/3}. \quad (1.96)$$

So for low frequency sensitivities the observing time significantly increases. For instance, let's now consider the observing time for a BNS system across different detectors. We suppose $m_1 \sim m_2 \sim 1.4M_\odot$ and $f_{\text{merge}} \sim 1.5\text{kHz}$

- LVK, $f_{\text{low}} \sim 20\text{Hz} \rightarrow \tau \sim 4\text{mins}$
- ET, $f_{\text{low}} \sim 1\text{Hz} \rightarrow \tau \sim 5\text{days}$

Neglecting the radiation of the Earth, overlapping and doppler effect. Similarly, the BBH are foreseen to be detectable in LISA band even of $\tau \sim 1$ month!

Another key aspect of the GWs is the relation between amplitude and source distance. By solving the linearized Einstein Equations, we find that the strain h scales as

$$h \sim \frac{1}{d} \left(\frac{GM_c}{c^3} \right)^{5/3} (\pi f)^{2/3} \quad (1.97)$$

Where d is the luminosity distance. If we consider the typical BNS at galactic distance $d \sim 400\text{Mpc}$ and $f \sim 1.4\text{kHz}$, this yields a strain amplitude on the order $h \sim 10^{-21}$. This amplitude-distance relation implies that, known the intrinsic parameters of the source, such as M_c , one can infer the distance d , allowing the use of GW as standard sirens in cosmology. The greater sensitivity of a detector, the farther the universe can be probed. Third-generation ground-based detectors such as the Einstein Telescope and Cosmic Explorer are expected to detect stellar-mass binary mergers out to redshifts of order $z \leq 10 - 20$. In addition, in scenarios where a significant population of primordial black holes exists, forecasts suggest that ET and CE could observe mergers at $z \geq 30$, well beyond the redshift range accessible to current detectors. For context, the first generation of stars (Population III) is expected to form at redshifts $z \sim 20 - 30$ [57–60]

Chapter 2

Detection of the inspiral binaries: instruments, techniques and pipelines

2.1 Interferometers

The development of the technologies necessary to measure GW events produced by the inspiral of compact binaries has required decades of sustained effort. Today, the detection of such events relies on interferometers (IFOs). One of the most historically significant scientific application of interferometry — aside from the detection of GWs — was the Michelson-Morley (MM) experiment [61], remarkably important for another Einstein’s theory of relativity, namely special relativity.

An IFO consists of two perpendicular arms of length L . A monochromatic laser beam is emitted from a source and directed along these arms. At the intersection point of the arms, a beam-splitter — i.e., a semi-transparent mirror — is placed to divide the incoming laser beam into two beams of equal probability amplitude, in accordance with the IFO geometry. Each of these beams travels along one of the arms, reflects off a mirror positioned at the end of the arm, and returns to the beam-splitter. Upon recombination at the beam splitter, the resulting interference pattern is measured by a photodetector. The intensity of light detected depends on variation in the optical path length L , since any change affects the relative phase of the two beams and thereby the interference pattern (see Fig. 2.1 for reference).

When a GW passes through the IFO, it induces a differential change in the length of the arms, modulating the interference pattern accordingly. Let us denote the angular frequency of the laser by ω_L , which is related to the laser’s wave number via $\kappa_L = \frac{\omega_L}{c}$. The spatial component of the laser’s electric field can be generically expressed as

$$E_0 e^{i\omega_L t + i\vec{\kappa}_L \cdot \vec{x}} \quad (2.1)$$

Let us consider a photon emitted by the laser source that reaches the beam splitter at $t = t_0$. The photon then propagates along each one of the IFO arms length L , denoted L_x for the arm aligned with

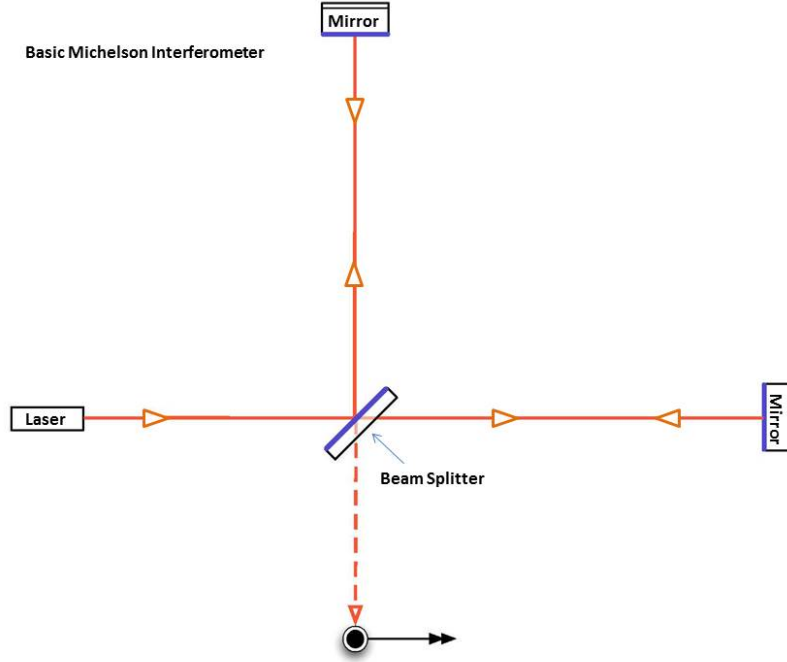


Figure 2.1 Example of Michelson-Morlet IFO - credits to [6]

the x -axis, and L_y for the arm aligned with the y -axis. Upon reaching the end of the arm, the photon is reflected by the mirror and returns to the beam splitter. The round-trip travel time is therefore $t = t_0 + \frac{2L_{x,y}}{c}$. At this time, the electric fields returning from both arms recombine at the beam splitter. These fields are given by:

$$E_x = -\frac{1}{2}E_0e^{-i\omega_L t + 2i\kappa_L L_x} \quad (2.2)$$

$$E_y = \frac{1}{2}E_0e^{-i\omega_L t + 2i\kappa_L L_y} \quad (2.3)$$

The minus sign in the expression for the x -component arises from the overall reflection effect at the beam splitter¹. The classical field energy measured by the photodetector is given by

$$P = |E_{out}|^2 = |E_x + E_y|^2 = E_0^2 \sin^2[\kappa_L(L_y - L_x)] \quad (2.4)$$

We observe that a variation in the arm lengths corresponds to variation in the measured power. Now, let us consider working in TT gauge. In this coordinates system, the effect of a passing GW does not manifest as a displacement of the mirrors' coordinates. Instead, the GW's influence is encoded in the time of flight of the photons, resulting in a variation in the phase $\Delta\phi_{x,y}$. For simplicity, we consider

¹see M. Maggiore 'Gravitational Waves' [21], Volume I, pp. 471

a linearly polarized GW with + polarization, propagating along the z -axis. In this case, the induced strain affects the x and y -components in the interferometer plane as follows:

$$h_+(t) = h_0 \cos(\Omega t) \quad (2.5)$$

Consider now the fact that a photon travels along null geodesics. In the TT gauge, the physical effect of a GW manifests through changes in the photon's time of flight, which translates into a variation in its phase $\Delta\phi_{x,y}$. This phase shift encodes the influence of the GW on the light propagation along the interferometer arms, and can be expressed as follows:

$$E^{(x,y)} = \mp \frac{1}{2} E_0 e^{-i\omega_L \left(t - \frac{L_{x,y}}{c} \right) + i\Delta\phi_{x,y}(t)} \quad (2.6)$$

$$\Delta\phi_{x,y} = \pm h_0 \frac{\omega_L L_{x,y}}{c} \text{sinc} \left(\Omega \frac{L_{x,y}}{c} \right) \cos \left[\Omega \left(t - \frac{L_{x,y}}{c} \right) \right] \quad (2.7)$$

The power measured by the photodiode is then given by

$$P \sim |E_{tot}|^2 = \frac{P_0}{2} \{1 - \cos[2\phi_0 + \Delta\phi_{MM}]\} \quad (2.8)$$

Where $\phi_0 = i\kappa_L(L_x + L_y)$ in the unperturbed phase, and $\Delta\phi_{MM} = 2\Delta\phi_x$ represents the differential phase shift induced by the GW's passage. In particular, in equations 2.7 we observe that the response of the interferometer exhibits a dependence on both the arm length L and the GW angular frequency Ω . In particular, the interferometer's sensitivity is maximized when condition $\frac{\Omega L}{c} = \frac{\pi}{2}$ is satisfied. This result is written in terms of gravitational wave frequency f_{GW} as

$$L = 750 \text{km} \left(\frac{100 \text{Hz}}{f_{GW}} \right) \quad (2.9)$$

Clearly, constructing interferometers of optimal length derived above is unfeasible on Earth. In practice, current instruments such as LIGO and Virgo have arm length of a few kilometers—specifically 4km for LIGO and 3km for Virgo. The detection of gravitational waves using such comparatively compact IFOs is made possible through the use of Fabry-Perot cavities. These optical cavities effectively increase the photon path length by causing the light to bounce multiple times between two highly reflective mirrors, thereby enhancing the interferometer's sensitivity to phase changes induced by GW. A key parameter governing the behavior of a Fabry-Perot cavity is the free spectral range, defined as $\Delta\omega_L = \frac{\pi c}{L}$, where L is the length of the cavity. This quantity determines the frequency spacing between resonant modes of the electromagnetic field within the cavity. Another important parameter is the full-width-half-maximum of the resonance, given by $\delta\omega_L = \frac{c}{L} \frac{1-r_1 r_2}{\sqrt{r_1 r_2}}$. $r_{1,2}$, where r_1 and r_2 denote the amplitude reflectivity coefficients of the two mirrors forming the cavity. Another important quantity associated to the Fabry-Perot cavity is the finesse defined as $\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1-r_1 r_2}$. The finesse characterizes the sharpness of the cavity resonances and is directly related to the photon




Figure 2.2 Example of MM IFO with FP cavity - credits to [6]

storage time, that is the average time a photon remains confined within the cavity. This storage time is defined as

$$\tau_s = \frac{2L}{c} \frac{1}{1-r_1^2} \Big|_{r_1 \rightarrow 1} \sim \frac{L}{c} \frac{\mathcal{F}}{\pi} \quad (2.10)$$

It can be shown that the use of Fabry-Perot cavities enhances the sensitivity to phase shift by a factor of approximately $\sim \frac{2}{\pi} \mathcal{F}$ [21]. This enhancement makes possible to achieve high sensitivity with significant more compact interferometer- such as those currently employed by the LVK collaboration.

2.2 Noise sources

The sensitivity required by gravitational wave interferometers is extraordinarily high. As previously discussed, the effect of a GW corresponds to a relative displacement of the interferometer mirrors at the order of $\sim 10^{-21}$. For an arm length of $L = 4\text{km}$, this implies an absolute variation $\Delta L \sim 10^{-18}\text{m}$; approximately one thousand times smaller than the diameter of a proton. Achieving such precision necessitates exceptional isolation from noise sources, which is critical for the instrument's performance. One of the most significant limitation arises from the quantum noise, which includes shot noise (affecting the measurement of the phase) and the radiation pressure (affecting the motion of the mirror). Together, these define the standard quantum limit. In addition to quantum noise, there is also

displacement noise, which receive contributions from seismic, Newtonian and thermal noise. Other sources of noise include contamination within the vacuum tube (e.g. residual gas molecules), scattered light and the fluctuation in laser power. this plurality of effect can compromise and deteriorate the IFO's sensibility. In the following section we examine some of these noise sources more in detail, focusing on their physical origin and the strategies employed to mitigate their impact.

2.2.1 Shot noise

Let us consider the laser beam as it reaches the photodetector. Due to the quantum nature of light, the laser can be described as a stream of discrete photons. Let N_γ denote the number of photons incident on the photodetector during a time interval T , with each photon carrying an energy $\hbar\omega_L$, where ω_L is the angular frequency of the laser. The average power measured over this time interval is therefore given by:

$$P = N_\gamma \hbar\omega_L \frac{1}{T} \quad (2.11)$$

Assuming that the photons arrives at the photodetector randomly, the detection process follows a Poissonian distribution. Consequently, the statistical fluctuation in the number of detected photons is $\Delta N = \sqrt{N_\gamma}$. This translates into a fluctuation in the measured power, known as shot noise, given by

$$(\Delta P)_{shot} = \left(\frac{\hbar\omega_L}{T} P \right)^{\frac{1}{2}} \quad (2.12)$$

Where P is the average power, T is the integration time, and ω_L is the laser angular frequency. We now compare this shot noise to the power fluctuation induced by the passage of a GW. As a first step, we consider the simplified case of an interferometer without Fabry-Perot cavity. From Eq. 2.8, the output power measured in the absence if a gravitational wave is:

$$P = P_0 \sin^2(\phi_0) \quad (2.13)$$

Where P_0 is the maximum power and ϕ_0 is the static phase difference between the interferometer arm. The corresponding fluctuation due to shot noise $(\Delta P)_{shot}$ is then

$$(\Delta P)_{shot} = \left(\frac{\hbar\omega_L}{T} P_0 \right)^{\frac{1}{2}} |\sin(\phi_0)| \quad (2.14)$$

When a gravitational wave passes through the interferometer, it modulates the optical path lengths and induces a time-dependent phase shift. The resulting fluctuation in the output power is given by:

$$(\Delta P)_{GW} = \frac{P_0}{2} |\sin(2\phi_0)| \Delta\phi_{MM} \quad (2.15)$$

Where $\Delta\phi_{MM}$ is the differential phase shift due to the GW. To evaluate $\Delta\phi_{MM}$, we assume, for simplicity, a linearly polarized GW + polarized, propagating along the \hat{z} direction. in the long-wavelength limit and $\frac{f_{GW}L}{c} \ll 1$, the phase shift is approximately

$$|\Delta\phi_{MM}| = \frac{4\pi L}{\lambda_L} h_0 \quad (2.16)$$

$$(\Delta P)_{GW} = \frac{P_0}{2} |\sin(2\phi_0)| \frac{4\pi L}{\lambda_L} h_0 \quad (2.17)$$

Where λ_L is the laser wavelength and h_0 is the GW strain amplitude. Comparing the Gravitational Wave induced power fluctuation $(\Delta P)_{GW}$ with the shot noise fluctuation $(\Delta P)_{shot}$ over an observation time T , we obtain

$$\frac{(\Delta P)_{GW}}{(\Delta P)_{shot}} = \left(\frac{P_0 T}{\hbar \omega_L} \right)^{\frac{1}{2}} \frac{4\pi L}{\lambda_L} h_0 |\cos(\phi_0)| \quad (2.18)$$

Evaluating this ratio at the nominal working point where $\cos(\phi_0) = \frac{1}{\sqrt{2}}$, we find

$$\frac{(\Delta P)_{GW}}{(\Delta P)_{shot}} = \left(\frac{P_0 T}{2\hbar \omega_L} \right)^{\frac{1}{2}} \frac{4\pi L}{\lambda_L} h_0 \quad (2.19)$$

Since the GW strain $h(t)$, as defined in 2.5, is periodic, we compare the time-averaged GW-induced fluctuation to the instrument's intrinsic fluctuation over the same T . It happens that a sinusoidal signal $h(t)$ is detectable when $(\Delta P)_{GW} = (\Delta P)_{shot}$ which yields the minimal detectable amplitude

$$h_{\min}(T) = \frac{\lambda_L}{4\pi L} \left(\frac{2\hbar \omega_L}{P_0 T} \right)^{\frac{1}{2}} \quad (2.20)$$

It is convenient to define an intrinsic (time-independent) strain sensitivity as amplitude spectral density (ASD) $S_n(f)^{\frac{1}{2}}$, obtained by multiplying h_{\min} by \sqrt{T} , as

$$S_n(f)^{\frac{1}{2}}|_{shot} = \frac{\lambda_L}{4\pi L} \left(\frac{2\hbar \omega_L}{P_0} \right)^{\frac{1}{2}} \quad (2.21)$$

This quantity has units $\text{Hz}^{-1/2}$ and characterizes the detector's shot-noise-limited response to strain, per $\sqrt{\text{Hz}}$, in the absence of any actual signal. With similar computation it is possible to obtain the strain sensitivity for a Fabry-Perot cavity, that results in

$$S_n(f)^{\frac{1}{2}}|_{shot} = \frac{1}{\mathcal{F}L} \left(\frac{4\pi \lambda_L c}{\eta P_{bs}} \right)^{\frac{1}{2}} \sqrt{1 + \left(\frac{f}{f_p} \right)^{\frac{1}{2}}} \quad (2.22)$$

Where P_{bs} is the power at the beam splitter after recycling, η is the efficiency of the photodiode and f_p is the pole frequency depending on the Fabry-Perot.

2.2.2 Radiation pressure

From Eq. 2.22, it is evident that increasing the laser power at the beam splitter P_{bs} reduces the contribution of shot noise to the strain sensitivity. This is expected, as a higher flux of photons decreases the relative uncertainty due to the quantum fluctuation in the photon number, which scales

as $\sqrt{N_\gamma}$. However, the interaction between the laser photons and the interferometer mirrors introduces an important counteracting effect. Gravitational waves induce a mirror displacement, but increasing the laser power - and hence the number of photons - amplifies the radiation pressure exerted on the mirrors. Each photons imparts a small moment upon reflection, resulting in a fluctuating radiation pressure force. So, we have to deal with a stochastic force that acts on our test masses. The strain sensitivity corresponding to this effect considering a Fabry-Perot cavity is

$$S_n(f)^{\frac{1}{2}}|_{RP} = \frac{16\sqrt{2}\mathcal{F}}{ML(2\pi f)^2} \sqrt{\frac{\hbar}{2\pi} \frac{P_{bs}}{\lambda_L c}} \frac{1}{\sqrt{1 + (f/f_p)^2}} \quad (2.23)$$

Where M is the mass of the mirror. This expression shows that, unlike shot noise, radiation pressure noise increases with laser power and dominates at low frequencies due to the f^{-2} dependence. At this point we can consider the total quantum-limited strain sensitivity as the sum of the contribution from the shot noise and radiation pressure noise:

$$S_n(f)|_{opt} = S_n(f)|_{shot} + S_n(f)|_{RP} \quad (2.24)$$

This quantity can be minimized with respect to a specific value of frequency f_0 corresponding to an optimal value of $S_n(f)$ that is the Standart Quantum Limit (SQL).

$$S_{SQL}^{\frac{1}{2}}(f) = \frac{1}{2\pi f L} \sqrt{\frac{8\hbar}{M}} \quad (2.25)$$

An example of the strain sensitivity due to the SQL is depicted in 2.4a.

2.2.3 Other noise sources

Beyond the quantum noise, several noise sources must be taken into account. One of the most important at low frequencies is the seismic noise, which originates from a variety of sources, ranging from persistent micro-seismic activity (ocean waves, micro-earthquakes) to human activity such as road traffic near the detector site. This noise scales as $S^{\frac{1}{2}}(f) \sim A \left(\frac{1\text{Hz}}{f^2}\right) \text{mHz}^{\frac{1}{2}}$, with $A \sim 10^{-7}$. At frequencies of a few hertz, this motion is many orders of magnitude larger than the mirror displacement caused by a gravitational wave. To suppress this disturbance, advanced interferometers employ seismic-isolation systems, known as superattenuators (Fig. 2.3), which provide extremely large attenuation at low frequencies. This structure consists of a cascade of pendulums, each with its own resonance frequency, providing passive isolation above resonance and active control below it. In Virgo, the seismic-isolation system is an $\sim 8\text{m}$ -tall structure composed of an inverted-pendulum stage at the top-used to provide horizontal isolation below 0.1Hz —followed by several pendulum stages and the marionette— mirrors suspension. The complete superattenuator chain provides a passive attenuation of $\sim 10^{12}$ a frequencies above 10Hz, reducing seismic motion to within the strain sensitivity required for gravitational-wave detection [7]. The overall design ensures that seismic noise




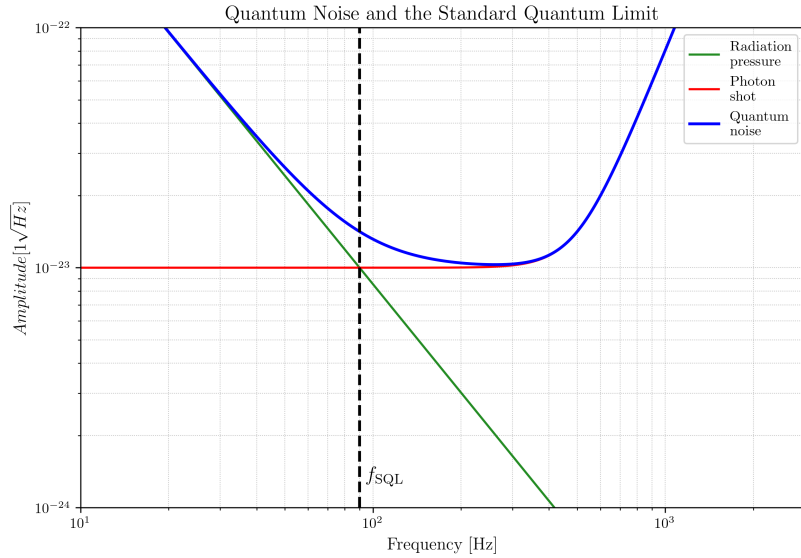
Figure 2.3 The scheme of the superattenuator. A Virgo mirror is supported in a vacuum by this kind of structure, which dampens seismic vibrations. It is a chain of pendula hanging from an upper platform and supported by three legs attached to ground, forming an inverted pendulum. — credits to [7]

is suppressed below the thermal-noise level above $\sim 4\text{Hz}$, enabling the interferometer to operate in astrophysical relevant band.

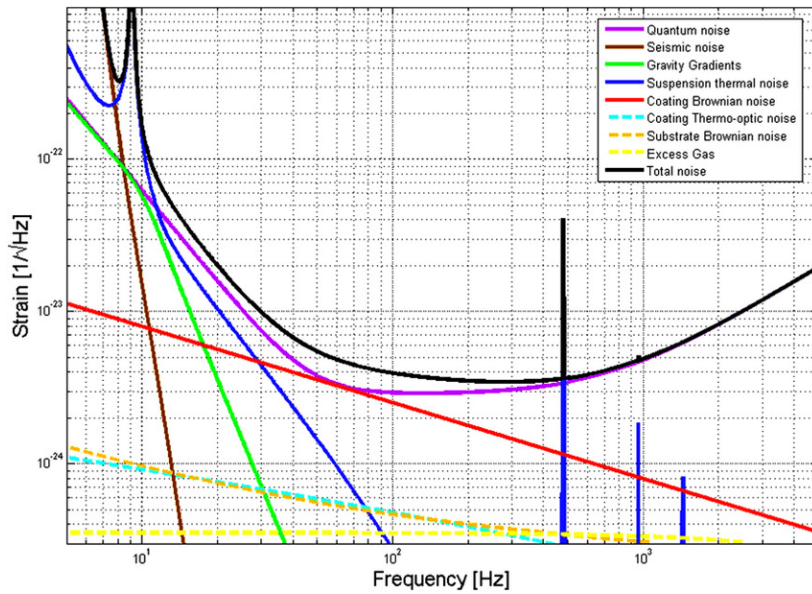
On the same nature, there is the Newtonian noise source that consists on the gradient of gravity field generated by several factors, such as the density of the soil where the IFO is build, or even pressure waves and temperature gradient [62]. This noise cannot be eliminated, since the gravitational force cannot be screened.

Thermal noise in the test-mass substrates, coatings, and suspensions is a fundamental limitation set by mechanical dissipation through the fluctuation-dissipation theorem [63]. In current detectors, the dominant contribution in the most sensitive band comes from the Brownian motion of the high-reflectivity mirror coatings [64–66]. This noise can be mitigated by employing low-loss substrate and coating materials, operating at cryogenic temperature, increasing the beam size on the mirrors, and using optimized multilayer coatings with reduced mechanical loss [67, 66, 68]. In addition to thermal noise, several technical noise sources must also be controlled. These include contamination by micro-particles inside the beam tube, scattered-light noise produced when stray light recombines with the main beam at the photodetector, and fluctuations in laser power. The combined effect of all

these contributions defines the interferometer’s noise budget, an example of which is illustrated in 2.4b.



(a) Example of Quantum Noise given the RP and SN contribution



(b) Noise Budget for ALIGO, credits to [69]

Figure 2.4 Comparison of quantum noise (Top) and ALIGO noise budget (Bottom).

2.3 Detecting the Gravitational Waves - Data analysis and techniques

In Eq. 2.21 we introduced the Strain Sensitivity by considering the ratio between the fluctuation to the input of photodetector caused by the passage of a GW signal and the noise caused by the shot

noise and taking out the GW amplitude h_0 and the time averaging T . We can introduce in a more formal-way this quantity, that we will refer to as Noise Spectral Density, that is fundamental for analysing the data stream provided by the IFOs. Suppose that the detector is a linear system, where at its input there is the signal we want to detect. We first recall that the output of a detector, that in this case is an IFO, it is a scalar number, while a GW that passes through this instrument is a tensor h_{ij} . This means that the output of the detector will be

$$h(t) = D^{ij}h_{ij} \quad (2.26)$$

where D^{ij} is the detector tensor and its form depends of the type of detector we are considering; this quantity also defines the detector pattern function, that is a quantity that incorporate the direction of propagation of the wave [70]. At this point, we can simply think that the detector output's is composed by the sum of $h(t)$ and a random noise $n(t)$ that gives

$$s(t) = h(t) + n(t) \quad (2.27)$$

The $n(t)$ is the detector's noise, where the sources are those described in the previous sections. For simplicity, we suppose now that the noise is stationary, so the different Fourier components are uncorrelated, that means that the average is²

$$\langle \tilde{n}^*(f)\tilde{n}(f') \rangle = \delta(f - f') \frac{1}{2} S_n(f) \quad (2.28)$$

Here we see a formal definition of the Noise Spectral Density firstly introduced in Eq. 2.21. From Eq. 2.28 we see that this quantity is even $S_n(-f) = S_n(f)$, since $n(t) \in \mathbb{R}$. We can also assume that $\langle n(t) \rangle = 0$. In the actual form, the left side of Eq. 2.28 diverges if $f = f'$, but in practical application we measure the strain, and so the noise, on a limited interval of time T , so if we consider the interval $\left[-\frac{T}{2}, \frac{T}{2}\right]$ then

$$\delta(f = 0) \rightarrow \left[\int_{-T/2}^{T/2} dt e^{2i\pi ft} \right] \Big|_{f=0} = T \quad (2.29)$$

So Eq. 2.28 at $f = f'$ is

$$\langle |\tilde{n}(f)|^2 \rangle = \frac{1}{2} S_n(f) T \quad (2.30)$$

Since the resolution in frequency is $\Delta f = \frac{1}{T}$, we can write Eq. 2.30 as

$$\frac{1}{s} S_n(f) = \langle |\tilde{n}(f)|^2 \rangle \Delta f \quad (2.31)$$

²Strictly speaking, the angle brackets denote an ensemble average over realizations of the process. In practice we approximate this by averaging over time, using segments of duration T that are separated by time-shifts longer than the noise correlation time, so that the segments are effectively uncorrelated (ergodic approximation).

inserting this expression, taking into account also the factor $1/2$, in Eq. 2.28, we obtain $\langle n^2(t) \rangle$ by integrating $S_n(f)$ over the physical frequencies $f \in [0, \infty)$

$$\langle n^2(f) \rangle = \int_0^\infty df S_n(f) \quad (2.32)$$

This quantity is fundamental since it quantifies, given a noise source, the sensitivity that the detector is capable of, and as we will see next, is a basic ingredient for a data analysis technique that is vastly used in this research field: the matched-filtering.

2.3.1 Matched-filtering

From Eq. 2.27, we see that a signal would be directly detectable only if $\frac{|h(t)|}{|n(t)|} \gg 1$, which is rather unrealistic, since in Earth-based detectors the signal is effectively buried in noise. A possible solution to this problem relies on the matched-filtering technique. The basic idea is that if the signal is known *a priori*, it can be extracted from the noise. Suppose that the detector output is given by Eq. 2.27, and that $h(t)$ is known. By multiplying the expression for the signal and integrating over an observation time T , we obtain the following equation

$$\frac{1}{T} \int_0^T dt s(t) h(t) = \frac{1}{T} \int_0^T dt h^2(t) + \frac{1}{T} \int_0^T dt n(t) h(t) \quad (2.33)$$

For large T , the terms on the right-hand side exhibit the following asymptotic behaviors:

$$\frac{1}{T} \int_0^T dt h^2(t) \sim h_0^2 \quad (2.34)$$

$$\frac{1}{T} \int_0^T dt n(t) h(t) \sim \left(\frac{\tau_0}{T}\right)^{\frac{1}{2}} n_0 h_0 \quad (2.35)$$

where h_0 and n_0 denote the characteristic amplitudes of $h(t)$ and $n(t)$, respectively, and τ_0 is the typical characteristic timescale of $h(t)$. As $T \rightarrow \infty$, the second term of Eq. 2.33 averages to zero. Notably, we do not require $h_0 > n_0$; it is sufficient that $h_0 > \left(\frac{\tau_0}{T}\right)^{\frac{1}{2}} n_0$. Let us formalize this result more rigorously. We define

$$\hat{s} = \int_{-\infty}^{\infty} dt s(t) K(t) \quad (2.36)$$

where $K(t)$ is the filter function. Next, we can introduce two additional quantities:

- S , the expectation value of \hat{s} when the signal is present, and
- N , the root-mean-square value of \hat{s} when the signal is absent.

Since we assumed $\langle n(t) \rangle = 0$, it follows that

$$S = \int_{-\infty}^{\infty} dt \langle s(t) \rangle K(t) = \int_{-\infty}^{\infty} dt h(t) K(t) = \int_{-\infty}^{\infty} df \tilde{h}(f) \tilde{K}^*(f) \quad (2.37)$$

Clearly, \tilde{h} and \tilde{K} are the Fourier transform, respectively, of h and K . The corresponding noise contribution is given by

$$N^2 = [\langle \hat{s}^2(t) \rangle - \langle \hat{s}(t) \rangle^2] = \int_{-\infty}^{\infty} df \frac{1}{2} S_n(f) |\tilde{K}(f)|^2 \quad (2.38)$$

Where $S_n(f)$ is the power spectral density of the noise. Combining these expressions, we define the signal-to-noise ratio (SNR), $\frac{S}{N}$ as

$$\frac{S}{N} = \frac{\int_{-\infty}^{\infty} df \tilde{h}(f) \tilde{K}^*(f)}{[\int_{-\infty}^{\infty} df \frac{1}{2} S_n(f) |\tilde{K}(f)|^2]^{1/2}} \quad (2.39)$$

It can be shown that the filter $K(t)$ that maximizes S/N is

$$\tilde{K}(f) = \frac{\tilde{h}(f)}{S_n(f)} \quad (2.40)$$

This equation defines the matched-filter. Another way to view this result is by defining the scalar product $\langle \cdot | \cdot \rangle$ between two real functions $A(t)$ and $B(t)$ as

$$\langle A | B \rangle = \text{Re} \int_{-\infty}^{\infty} df \frac{\tilde{A}^*(f) \tilde{B}(f)}{(1/2) S_n(f)} = 4 \text{Re} \int_0^{\infty} df \frac{\tilde{A}^*(f) \tilde{B}(f)}{S_n(f)} \quad (2.41)$$

Imposing $A = h$ and $B = n$, by using Eq. 2.39 and 2.40 we find

$$\left(\frac{S}{N}\right)^2 = \langle h | h \rangle = 4 \int_0^{\infty} df \frac{|\tilde{h}(f)|^2}{S_n(f)} \quad (2.42)$$

This expression is the optimal signal-to-noise ratio achievable through matched-filtering.

2.3.2 The matched-filtering statistics

At this stage, we have a tool that allows us to extract a GW signal buried in noise. The price we pay for this capability is the requirement of a set of possible signals, consistent with the expected astrophysical population of interest. In practice, this means we need physical models that specifies the parameters of the systems generating the GWs, together with the corresponding solutions of the Einstein Equation for these systems. Once this set of possible waveforms — known as template bank — is computed, the standard analysis consists of performing matched-filtering of the data using the waveforms in the template bank. This procedure gives a list of candidate events, each associated with a signal-to-noise ratio. A natural question arises: what is the statistical significance of an event with a given SNR? The general idea is to impose an SNR threshold such that most noise triggers fall below it, while true GW signals exceed it. The simplest type of noise encountered in the search is the Gaussian noise (GN). In this context we define

$$\rho = \frac{\hat{s}}{N} \quad (2.43)$$

Where \hat{s} is defined in Eq. 2.36 and N is the square root of Eq. 2.38. ρ correspond to the SNR, with its expectation value given by $\langle \rho \rangle = \frac{\hat{s}}{N}$. We recall from Eq. 2.36 that

$$\hat{s} = \int_{-\infty}^{\infty} dt s(t)K(t) = \int_{-\infty}^{\infty} dt [h(t) + n(t)]K(t) \quad (2.44)$$

In the absence of a signal, \hat{s} reduces to a random variable determined solely by the noise $n(t)$, which has zero mean and unit variance, taking into account the normalization. Therefore, the probability density of ρ in absence of a signal is

$$p(\rho|h=0)d\rho = \frac{1}{\sqrt{2\pi}} e^{-\frac{\rho^2}{2}} d\rho \quad (2.45)$$

Conversely, suppose a signal is present. Then

$$\rho = \frac{\hat{s}}{N} = \frac{\int_{-\infty}^{\infty} dt h(t)}{N} + \frac{\int_{-\infty}^{\infty} dt n(t)}{N} = \bar{\rho} + \hat{n} \quad (2.46)$$

where $\bar{\rho}$ is the SNR output from matched-filtering in the presence of a signal, and \hat{n} represents the normalized noise contribution. Rewriting this relation in terms of \hat{n} , we find that $\rho - \bar{\rho}$ is a Gaussian random variable with zero mean and unit variance. It follows that the probability distribution of ρ when a signal is with true SNR $\bar{\rho}$ is present is

$$p(\rho|\bar{\rho})d\rho = \frac{1}{\sqrt{2\pi}} e^{-\frac{(\rho-\bar{\rho})^2}{2}} d\rho \quad (2.47)$$

By defining $R = \rho^2$, that is the SNR in energy, we find that the probability distribution for R when a signal is present and shows a SNR in energy as $\bar{R} = \bar{\rho}^2$ results in

$$P(R|\bar{R}) = \frac{1}{\sqrt{2\pi R}} e^{-\frac{\bar{R}+R}{2}} \cosh[\sqrt{R\bar{R}}] dR \quad (2.48)$$

In Fig. 2.5a is depicted the different distribution for R when a signal is absent $\bar{R} = 0$ or present, with for instance $\bar{R} = 10$. The probability distribution behaves very differently depending on whether a signal is present or not, as shown in the picture. In the absence of a signal, the shape is significantly different from the case where a true signal is present. This result suggests that we can define a false-alarm-probability by setting a threshold R_t above which noise-only triggers are unlikely to appear. The false-alarm-probability is given by

$$P_{FA} = \int_{R_t}^{\infty} dR P(R|\bar{R}=0) = 2 \int_{\rho_t}^{\infty} d\rho e^{-\frac{\rho^2}{2}} \quad (2.49)$$

The threshold R_t can thus be chosen to achieve a desired maximum false-alarm rate that we are willing to tolerate. Now, suppose we have two outputs x and y , each with its own Gaussian noise, and we combine them in quadrature such that $\rho^2 = x^2 + y^2$. This situation is often present since usually we matched-filter the data against unknown signal, therefore the signal must be maximized with respect

to in-phase and in-quadrature outputs, that correspond in this case to x and y . Similarly to what we have just seen, the probability distribution of ρ in absence of signal is then

$$p(\rho|h=0) = \rho e^{-\frac{\rho^2}{2}} \quad (2.50)$$

where we have moved to polar coordinates (ρ, θ) and integrated over the phase θ . This corresponds to a Rayleigh distribution or a χ^2 distribution with two degrees of freedom. Now consider the case where a GW signal event is present, so the true signal contributes \bar{x} and \bar{y} . In this case, the joint distribution for x and y is

$$p(x, y|\bar{x}, \bar{y}) = \frac{1}{2\pi} e^{-\frac{1}{2}[(x-\bar{x})^2 + (y-\bar{y})^2]} \quad (2.51)$$

Switching to polar coordinates and defining $R = \rho^2$ and $\bar{R} = \bar{x}^2 + \bar{y}^2$, it can be shown that resulting distribution is

$$P(R|\bar{R}) = \frac{1}{2} e^{-\frac{R+\bar{R}}{2}} I_0(\sqrt{R\bar{R}}) \quad (2.52)$$

where I_0 is the modified Bessel function, which arises from the integration over the phase θ . This distribution corresponds to a non-central χ^2 distribution with 2 degrees of freedom, Fig. 2.5b shows the distribution of R sampled from a SNR times series when a signal is present in the data. The mean and variance of such distribution are given by


$$\langle R \rangle = 2 + \bar{R} \quad (2.53)$$

$$\langle R^2 \rangle - \langle R \rangle^2 = 4 + 4\bar{R} \quad (2.54)$$


2.3.3 Coalescence of compact binaries

In this thesis we will analyze the gravitational-wave emission of compact binary coalescence systems, such as Binary Black-Holes (BBH), Binary Neutron-Star (BNS) or Neutron-Star Black-Hole (NSBH). In the first chapter we saw that those systems release a huge amount of energy in their final inspiral phase, and also their dynamics generate waveforms that can consist in many cycles that can be tracked for a long time by the Earth IFOs. In addition, solution to the Einstein Equations associated to these phenomena are given, that means that it is possible to build the templates needed for the matched-filtering analysis.

One point that must be stressed is that these waveform are described by a consistent number of parameters, some concerning the position in the sky or the arrival time, or some others such as the masses and the spins of the binary. In total, such systems show up to 15 parameters, implying a very large number of templates to search over. However, the dimensionality can be reduced by exploiting parameter dependencies and analytic maximizations [21].



(a) Probability distribution of $P(R|\bar{R})$ in the absence of a signal ($\bar{R} = 0$) and in the presence of a signal with $\bar{R} = 10$



(b) Simulated distribution of SNR for data containing an injected signal.

Figure 2.5

Among the extrinsic parameters, an important role is played by the arrival time t_* of the signal. The maximization over t_* can be efficiently performed by evaluating the matched filter as a function of time through a single inverse Fourier transform, and identifying the time at which the signal-to-noise ratio (SNR) is maximum.

A waveform with arrival time t_* can be written as a time translation of a reference waveform,

$$h(t; \theta, t_*) = h(t - t_*; \theta), \quad (2.55)$$

where θ denotes the intrinsic parameters. In the frequency domain this corresponds to

$$\tilde{h}(f; \theta, t_*) = \tilde{h}(f; \theta) e^{2\pi i f t_*}. \quad (2.56)$$

Substituting this expression into the matched-filter inner product (Eq. 2.41) shows that the dependence on t_* appears as a complex phase factor. The inverse Fourier transform of the frequency-domain integrand therefore yields the SNR time series, whose maximum directly determines the optimal arrival time t_* .

Two other parameters that can be eliminated analytically from the matched-filtering search are the overall amplitude A of the signal and the phase ϕ [21].

The optimal filter is defined modulo an arbitrary normalization, so that the absolute value of the amplitude A does not affect the search for the template that maximizes the signal-to-noise ratio. Maximizing the likelihood with respect to A yields a statistic proportional to

$$\frac{\langle h|s \rangle^2}{\langle h|h \rangle} \quad (2.57)$$

The phase ϕ can also be maximized analytically by writing the template as

$$h(t; \phi) = h_c(t) \cos \phi + h_s(t) \sin \phi \quad (2.58)$$

and introducing two new templates obtained as linear combinations of h_c and h_s ,

$$h_p = h_c \cos \phi_p + h_s \sin \phi_p \quad (2.59)$$

$$h_q = h_c \cos \phi_q + h_s \sin \phi_q \quad (2.60)$$

where the angles ϕ_p and ϕ_q are chosen such that $\langle h_p|h_q \rangle = 0$. In terms of these orthogonal templates, the statistic maximized over ϕ becomes a sum in quadrature of two matched-filter outputs,

$$\frac{\langle h|s \rangle^2}{\langle h|h \rangle} \longrightarrow \frac{\langle h_p|s \rangle^2}{\langle h_p|h_p \rangle} + \frac{\langle h_q|s \rangle^2}{\langle h_q|h_q \rangle} \quad (2.61)$$

so that maximizing over the phase is equivalent to maximizing the sum in quadrature of the outputs of two matched filters.

The matched-filtering output defines the signal-to-noise ratio time series $\rho(t)$

$$\rho(t) = 4 \left| \int_0^\infty \frac{\tilde{h}^*(f) \tilde{s}(t)}{S_n(f)} e^{2i\pi f t} df \right| \quad (2.62)$$

and an example of this observable is shown in Fig. 2.6, where we matched-filter the injected signal with a template sharing the same parameters, obtaining what is known as optimal SNR.




Figure 2.6 Matched-Filtering time series for an injected BBH of $5 - 5M_{\odot}$ at $800Mpc$ in Gaussian noise colored by the PSD of LIGO IFO. For simulations we use the Advanced LIGO zero-detuned high-power design sensitivity curve (aLIGOZeroDetHighPower) as implemented in LALSimulation [8].

2.3.4 Mismatch

Of course, the parameters' space that describes the waveform $h(t)$ is continuous, that means that in order to obtain the optimal filter for each astrophysical event, we would need an infinite number of waveforms, that is computationally unfeasible. To overcome this problem, we must rely on a limited number of possible signals to use for computing the matched-filtering. The set of this discrete signal is known as the template bank. The immediate consequence is that the sparsity of the templates will cause a loss in the SNR and, consequently, in the event rate, since some events will fall below the SNR's threshold due to the mismatch. In this sense, we must looking for a set of templates that is computationally acceptable and bearable in SNR's loss. We can describe a waveform templates by a function $T(t; \vec{\mu}, \vec{\lambda})$, whith $\vec{\lambda}$ are the intrinsic parameters, while $\vec{\mu}$ are the extrinsic one [71]. We can assume the templates T as normalized, so by Eq. 2.39 we have $\langle T(\vec{\mu}, \vec{\lambda}) | T(\vec{\mu}, \vec{\lambda}) \rangle = 1, \forall \vec{\mu}, \vec{\lambda}$. We already saw that the extrinsic parameters can be handled easier with respect to the intrinsic one, and those are the ones we will focus on. We can introduce the quantity named match M between two templates $\tilde{T}(f, \vec{\lambda})$ and $\tilde{T}(f, \vec{\lambda} + \Delta\vec{\lambda})$ that is

$$M(\vec{\lambda}, \Delta\vec{\lambda}) = \langle \tilde{T}(\vec{\lambda}) | \tilde{T}(\vec{\lambda} + \Delta\vec{\lambda}) \rangle \quad (2.63)$$

This quantity identifies the fraction of SNR obtained using a template T that differs in the intrinsic parameter space by $\Delta\vec{\lambda}$. In this way, it is possible to tell how close to each others are two templates. Clearly, if $\Delta\vec{\lambda} = 0$ then we obtain a maximum value of match $M = 1$. Around this value, we can perform an expansion in series at second order and obtain

$$M(\vec{\lambda}, \Delta\vec{\lambda}) \sim 1 + \frac{1}{2} \left(\frac{\partial^2 M}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right)_{\Delta\vec{\lambda}=0} \Delta\lambda^i \Delta\lambda^j \quad (2.64)$$

In this way we can introduce the metric

$$g_{ij}(\vec{\lambda}) = -\frac{1}{2} \left(\frac{\partial^2 M}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right)_{\Delta\vec{\lambda}=0} \quad (2.65)$$

that quantity naturally describes the mismatch between two nearby templates

$$1 - M = g_{ij} \Delta\lambda^i \Delta\lambda^j. \quad (2.66)$$

This quantity can be interpreted as a distance between two templates separated in intrinsic parameter space by $\Delta\lambda$. Then, we can define then the Minimal Match MM — also known as Fitting Factor or coverage — as the match between the signal and the nearest templates in the worst possible case. Given a set of template, two nearest templates is a couple where such distance is the smallest in the whole set. Consequently, the MM can be identified as the threshold on the quantity of SNR that can be loss by a bank, and so increasing this will require a thicker template banks, while relaxing this will cause a sparsity of the templates in the parameters space.

2.4 The Multi-Band Template Analysis Pipeline

At this point we can introduce a possible implementation of the tools described. We saw that in the context of compact binary coalescence (CBC), the data stream provided by the IFOs is analyzed using the matched-filtering technique. This method allows us in measuring events that are buried under the noise given a waveform that is 'similar' to the signal present in the data. In this section we will quantify the meaning of 'similar' and describe one of the actual algorithm running in the LVK collaboration for detecting the CBC events: the Multi-Band Template Analysis pipeline (MBTA) [72] [73] [74].

MBTA is a pipeline designed to search for CBC events, with a focus on online analysis. It has been used since S6/VR3 joint LIGO-Virgo run [75]. This algorithm has been developed to reduce the computational cost of the search by running matched-filtering analysis in multiple frequency bands and recombining coherently their output to recover the same signal of a regular search. MBTA has all the standard required for a CBC search: it is able to trigger candidates, It can verify the coincidence between detectors, compute both false alarm rate and probability of astrophysical origin (p_{astro}), which will be explained in detail in Sec. 2.4.2. and, finally, submit to Gravitational Waves Candidate Event Database (GraceDB). GraceDB is a service operated by the LIGO Scientific Collaboration. It provides a centralized location for aggregating and retrieving information about candidate gravitational-wave events. GraceDB provides an API for programmatic access, and a client package is available for interacting with the API [76].

Since MBTA is looking for CBC events, a fixed template bank is used with a typical minimal match of 0.97 to scan the parameters' space. As the pipeline's name suggests, the search is performed by splitting the matched-filtering procedure among two, or more, frequency bands. In this way a smaller template bank can be used, since fewer cycles of the signals must be tracked, and a reduced sampling rate can be used for the lower frequency band reducing the computational cost of the Fourier Transform needed for the matched-filtering.³ The full band signal is obtained by summing the output of the two frequency bands matched-filtering. Defining as $\tilde{h}(f)$ the output in the frequency space of the IFOs and $\tilde{T}(f, \vec{\lambda})$ a generic template from the bank we have

$$out(t, \vec{\lambda}) = \int_{f_{min}}^{f_{max}} \tilde{h}(f) \tilde{T}(f, \vec{\lambda}) df = \int_{f_{min}}^{f_c} \tilde{h}(f) \tilde{T}(f, \vec{\lambda}) df + \int_{f_c}^{f_{max}} \tilde{h}(f) \tilde{T}(f, \vec{\lambda}) df \quad (2.67)$$

The f_c is usually chosen in order to share the same amount of SNR among the frequency bands, and usually is $\sim 80Hz$.

For each frequency band, the matched filtering is computed with template banks independent from each others. The banks associated to these frequency bands are called Real Template (RT) banks, while the bank that contains the templates related to the whole frequency band is known as Virtual Template (VT). Each of this bank is required to have the same MM , for instance 0.97, and then the templates are one-to-one mapped between the RTs banks and the VT. The latter is the one that contains the parameters associated to the MBTA's trigger. The algorithm performs the matched-filtering time series in each frequency band both in phase and quadrature. So we have the following quantities

$$\langle h, RT_{LF} \rangle_P(t): \text{phase output of low frequency real template} \quad (2.68)$$

$$\langle h, RT_{LF} \rangle_Q(t): \text{quadrature output of low frequency real template} \quad (2.69)$$

$$\langle h, RT_{HF} \rangle_P(t): \text{phase output of high frequency real template} \quad (2.70)$$

$$\langle h, RT_{HF} \rangle_Q(t): \text{quadrature output of high frequency real template} \quad (2.71)$$

The operation in phase and quadrature are defined through eq. (2.39) and considering a phase term

$$\langle h, T \rangle_P = 4 \left| \int_0^\infty df \frac{\tilde{h}(f) \tilde{T}^*(f)}{S_n(f)} e^{2i\pi ft} \right| \quad (2.72)$$

$$\langle h, T \rangle_Q = 4 \left| \int_0^\infty df \frac{\tilde{h}(f) \tilde{T}^*(f)}{S_n(f)} e^{2i\pi ft - i\pi/2} \right| \quad (2.73)$$

³In addition, MBTA's carbon footprint is at low impact, since it requires $\sim 2\%$ of computing resource allocated online analyses at the main LVK low-latency computing centers [74].

Those time series must be rotated and aligned in time and phase to be coherently summed, it results then in

$$\langle h, VT \rangle_P(t) = \langle h, RT_{LF} \rangle_P(t) + \langle h, RT_{HF} \rangle_P(t)(t + \Delta t) \cos(\Delta\phi) - \langle h, RT_{HF} \rangle_Q(t)(t + \Delta t) \sin(\Delta\phi) \quad (2.74)$$

$$\langle h, VT \rangle_Q(t) = \langle h, RT_{LF} \rangle_Q(t) + \langle h, RT_{HF} \rangle_P(t)(t + \Delta t) \sin(\Delta\phi) + \langle h, RT_{HF} \rangle_Q(t)(t + \Delta t) \cos(\Delta\phi) \quad (2.75)$$

The offset Δt , $\Delta\phi$ are the time needed for the signal to chirp from the minimum frequency f_{min} to the cut-off frequency f_c and the phase at that point.

2.4.1 Single triggers and coincidence

Once the data stream provided by the IFOs is calibrated and conditioned with data quality informations [77], MBTA performs the matched-filtering analysis for each template in each frequency band as described above. For each unit of stretch processed data, the maximum value SNR (ρ), is compared to the minimum threshold ρ_{th} and the trigger is accepted if $\rho > \rho_{th}$. In order to limitate a glitch-dominated ρ distribution, a basic acceptance criterion is applied using a cut on the 2 degrees-of-freedom χ^2 observable which quantifies the agreement between the observed distribution of SNR across frequency and that expected from an astrophysical signal [78].

In the two-band implementation adopted here, the matched-filter output is split into two frequency bands, and the corresponding χ^2 statistic is computed from the difference between the SNR contributions in the two bands. A trigger is accepted if the χ^2 value satisfies

$$\chi^2 < \alpha (2 + \beta \rho^2), \quad (2.76)$$

where the constant term reflects the expected mean value of the χ^2 distribution for two degrees of freedom. The parameters α and β are empirically tuned and control the strictness of the veto: α sets the overall tolerance of the test, while β allows the threshold to increase with SNR, preventing the rejection of high-SNR signals that may exhibit small modeling inaccuracies. In this analysis we adopt $\alpha = 3$ and $\beta = 0.025$, which provide an efficient compromise between glitch rejection and signal retention.

In this direction it is possible to weight the ρ distribution [79] in order to penalize glitchy-triggers behavior pushing the corresponding statistics to lower statistical significance values. Following [80] we define the autocorrelation-based least-squares ξ_{PQ}^2 as

$$\xi_{PQ}^2 = \frac{1}{2\Delta t} \int_{t_0 - \frac{\Delta t}{2}}^{t_0 + \frac{\Delta t}{2}} \left\| \begin{pmatrix} \rho_P(t) \\ \rho_Q(t) \end{pmatrix} - \rho \mathcal{R} \begin{pmatrix} A_P(t - t_0) \\ A_Q(t - t_0) \end{pmatrix} \right\|^2 \quad (2.77)$$

Given a trigger measured at time t_0 , the corresponding complex SNR modulus and phase are ρ and ϕ . The rotation matrix \mathcal{R} is associated to the phase ϕ , the in-phase and in-quadrature matched-filter outputs ρ_P ρ_Q , and the autocorrelations of the template with its in-phase and in-quadrature A_P , A_Q . This quantity ξ_{PQ}^2 is then combined to the ρ statistics in such way

$$\begin{cases} \rho_{rw} = \rho & \text{if } \xi_{PQ}^2 \leq 1 \\ \rho_{rw} = \rho \left(\frac{A + (\xi_{PQ}^2)^\alpha}{A+1} \right)^{-1/\beta} & \text{if } \xi_{PQ}^2 > 1 \end{cases} \quad (2.78)$$

with $\alpha = 5$, $A = 10$ and $\beta = 8$. A further re-weight of the ρ_{rw} can be done taking into account the excess rate of triggers $R_{(\rho \geq \rho_{min})}$ compared to the rate of surviving triggers once the χ^2 cut is applied and the ρ_{rw} is computed $R_{(\rho_{rw} \geq \rho_{min})}$. In this way we obtain a new observable that quantifies the evidence of poor data quality. This rate is defined as

$$E_R(t_0) = \text{median}_{[t_0 + \delta t - 10s, t_0 + \delta t]} \delta R(t) \quad (2.79)$$

$$\delta R(t) = \frac{R_{(\rho \geq \rho_{min})} - R_{(\rho_{rw} \geq \rho_{min})}}{R_{(\rho \geq \rho_{min})}} \quad (2.80)$$

δR is sampled at 8Hz and E_R at 1Hz. To obtain the new statistics we now can do

$$\begin{cases} \rho_{rw, E_R} = \rho_{rw} & \text{if } E_R \leq 0.3 \\ \rho_{rw, E_R} = \rho_{rw} [1 - A(E_R - 0.3)^\alpha] & \text{if } E_r > 0.3 \end{cases} \quad (2.81)$$

The parameters' values are $A = 1$, $\alpha = 2$, $\delta t = 0s$ for online and $\delta t = 7s$ for offline. Recently, we decided to switch to a new possible approach that is replacing the E_R with the SNR-Excess. This observable is build by considering the maximum ρ time series over the template bank considering a time-window. This time series is then compared using a χ^2 statistics of a model based on a large population of injected signal in real data. What it is expected is a different behavior for noise and signals, with many templates will trigger for glitches while many less for signals. The new ranking statistics RS is

$$\begin{cases} RS = RS & \text{if } \chi_{SNR-Excess}^2 < 1 \\ RS = RS \left(\frac{1}{\chi_{SNR-Excess}^2} \right) & \text{otherwise} \end{cases} \quad (2.82)$$

The triggers that survive the cuts are defined as single triggers. If two or more triggers have been found by the same templates in two or more detectors, and in a time window physically compatible to the time of flight of the GW between the detectors, then a coincidence trigger is claimed. The time window considered for double coincidence as Hanford-Livingston (HL) is $15ms$, while for Hanford-Virgo (HV) or Livingston-Virgo (LV) is $35ms$. Furthermore, triple coincidence can be obtained by searching for HL and HV events that share the same H single-detector trigger. If an astrophysical

signal is measured by the IFOs, it is common that several templates will trigger within a close range of time, usually $20ms$. Since we need to choose only one possible trigger candidate among these, it is chosen the one with highest ranking statistics, using a procedure named clustering. Once coincidence triggers have been clustered, the ranking statistics of the single IFOs RS_i are summed in quadrature obtaining the combined ranking statistics cRS and finally a threshold is applied.

$$cRS_{ij}^2 = RS_i^2 + RS_j^2 + 2\ln(P_{ij}) \quad (2.83)$$

We have to highlight that a correction value P_{ij} is taken into account. This quantity depends on the arrival time and phase difference between the detectors, and it compared to the expected astrophysical distribution assuming uniform population [74].

2.4.2 False Alarm Rate and probability of being astrophysical source

Once a coincidence trigger with a given ranking-statistics is identified, we want to assign a confidence metric that quantifies whether it is a noise trigger or the probability that it is an astrophysical event. To do so, two related quantities have been defined. On the one hand, the False Alarm Rate (FAR) specifies how often noise trigger with a given ranking-statistics occur; on the other, the p_{astro} quantifies the probability that a trigger is of astrophysical origin [81] [82].

In O4a, the procedure for estimating these quantities consists in three steps. First, the FAR is computed at the candidate cRS ; second, the p_{astro} is obtained for that specific candidate and, finally, a mapping between p_{astro} and FAR is established.

For detector pair $i - j$ the FAR at a coincidence ranking-statistics threshold $\rho_{cRS,ij}$ is modeled as

$$FAR_{ij}(\rho_{cRS,ij}) = \frac{N_{ij}(\rho_{cRS,ij}) \cdot w_{ij}}{T_i T_j} \quad (2.84)$$

where $N_{ij}(\rho_{cRS,ij})$ is the number of background (fake) coincidences with ranking-statistics at or above the threshold $\rho_{cRS,ij}$, estimated empirically from time-shifted (non-physical) combinations of single-detector trigger lists under the assumption of independent noise in each detector. Here w_{ij} is the coincidence time window between detectors i and j and T_i, T_j are the corresponding analyzed lifetime duration. The FAR is expressed in yr^{-1} and represent the expected rate of noise triggers at the specified value of cRS .

On the other side, the p_{astro} probability can be defined as the ratio between the rate density of signals $R(k)_S$ and the rate of all the triggers detected, that corresponds to the sum $R(k)_S + R(k)_N$, given the latter as the rate density of noise triggers. Formally p_{astro} is

$$p_{\text{astro}} = \frac{R(k)_S}{R(k)_S + R(k)_N} \quad (2.85)$$

In general, these rate densities vary as a function of the ranking statistics k , which for MBTA correspond to cRS . The p_{astro} for MBTA is currently implemented as described in [83],

$$p_{\text{astro}}(cRS) = \frac{n_s(cRS)}{n_s(cRS) + n_b(cRS)} \quad (2.86)$$

where n_s and n_b are the densities of signals and background triggers as function of the cRS . Since these populations vary in the parameters space, this expression is build considering several bins that span among the chirp mass, the mass ratio and the possible detector combinations. In general these populations can be factorize in this way

$$n_s(cRS) = p(cRS|s)\Lambda_1 \quad (2.87)$$

$$n_b(cRS) = p(cRS|b)\Lambda_0 \quad (2.88)$$

The values Λ_1 and Λ_0 are the priors for the signal and noise events, that is the expected number, or rate, of signal or noise triggers that we expect to measure during a data acquisition campaign. The functions $p(cRS|s)$ and $p(cRS|b)$ quantify the distribution of, respectively, signal and noise among the statistics cRS . The background models are build taking into account the distribution of noise events in the possible parameters' space and coincidence combinations. On the other side, the signal distribution is obtained considering injections which have the astrophysical distribution foreseen by the models and taking into account a scaling factor that normalizes the number of injections with the expected number of detections, as described in details here [83]. In this way a relationship between p_{astro} and the ranking statistics cRS subsists via parametrization. For each trigger we measure the p_{astro} corresponding to the cRS in that parameters' space. In addition, the p_{astro} can be build as sum of possible astrophysical sources $p_{\text{astro}} = p_{\text{BNS}} + p_{\text{NSBH}} + p_{\text{BBH}}$, where these are obtained by multiplying p_{astro} for the corresponding source-wise relative rate in the specific bin.

Exploiting the p_{astro} , we developed a computation of the false alarm rate FAR as function of this metric. We can build this significance by estimating the number of background events above any given p_{astro} . The idea is to build a $FAR_{id}(p_{\text{astro}})$ taking into account the parameters' bin i and the detector combination d . This is done exploiting the relationships between p_{astro} and the cRS and the distribution of signal $n_s(cRS)$. Then we combine all the results from different detectors combination and obtain a single parametrization for the parameters' space bins

$$FAR_i(p_{\text{astro}}) = \sum_{d \in \{\text{All IFOs comb.}\}} w_d FAR_{(i,d)}(p_{\text{astro}}) \quad (2.89)$$

Where w_d is the fraction of observing time for each detector combination over the total. At this point, summing over the whole set of bins it is possible to obtain

$$FAR(p_{\text{astro}}) = \sum_i FAR_i(p_{\text{astro}}) \quad (2.90)$$

Chapter 3

Template bank construction

In the previous chapter we exposed the compact binary coalescence search analysis needed to detect a GW signal. One of the fundamental point that arose was building a template bank ensuring an acceptable coverage of the parameters space and a computational cost that depends on the number of templated used in matched-filtering. The discrete nature of the banks implies that if the signal does not correspond exactly to the template it is matched-filtered against, it results in a loss of SNR. The templates then are placed in order to guarantee a maximum loss in SNR around 3% in the whole parameters' space, taking into account the strain sensitivity curve of the IFOs (PSD).

The template bank MBTA deploys a 4-dimensions set build up by considering the masses m_1, m_2 and the dimensionless spin χ_1^z, χ_2^z assumed to be parallel (aligned or anti-aligned) with the angular momentum L . The choice of parallel spins is well justified: This restriction to parallel spins is well motivated: although generic spin misalignment induces orbital precession, several studies have shown that precessional effects have a limited impact on the detection efficiency of matched-filter searches based on aligned-spin templates [84, 85]. In particular, the loss in matched-filter signal-to-noise ratio due to neglecting precession is typically below $\sim 5\%$ for the majority of astrophysically relevant binary configurations, and rarely exceeds $\sim 10\%$ except for systems with large in-plane spins and unfavorable orientations [86, 87].

Consequently, restricting the search to aligned-spin templates does not result in a significant degradation in signal recovery, while it substantially reduces the dimensionality of the parameter space and the associated computational cost of template-bank construction.

Since MBTA performs matched-filtering in multiple frequency bands, it requires the construction of three distinct template banks. The first is the full-bandwidth bank that is referred to as Virtual Template Bank (VT). The other two are Real Template banks, one for low-frequency band (RT-LF) and one for the high frequency band (RT-HF). The RT banks are the ones actually used to perform the matched-filtering in their respective frequency ranges, while the VT bank represents the full-band equivalent templates and is used to record the effective triggers. Each virtual template corresponds uniquely to a pair of real templates, ensuring a one-to-one mapping between VT and RT banks. The following section briefly summarizes the production used to generate such banks.

3.1 Geometric, stochastic and hybrid placement

There are several ways to build a template bank, each one relies on the optimization problem of maximizing the coverage by reducing the number of templates. In Eq. 2.65 we saw that the metric, depending on the parameter space, is not generally flat. Yet it is possible to approximate it as locally flat, especially for low masses systems like BNS. This solution is implemented in `pycbc_geom_aligned_bank` [88]. In this way, one can analytically compute the metric and place the templates in such way that the coverage is guaranteed. This procedure is named *geometric placement*. The great advantage of this approach is that the computational cost is limited — since the metric is known — and the coverage is guaranteed. Unfortunately, it is not always possible to apply this solution, especially for high masses parameters’ space, where the metric curvature is much bigger: that’s why stochastic methods have been introduced.

The stochastic placement methods, as the name suggests, do not rely on the explicit, or approximated, expression of the metric. Instead, these approaches in general create several random points in the parameters space, and compute the overlap between the corresponding waveform. If there is an ‘overfit’ of the points, that correspond in one or more waveforms that match between each other above the required minimal-match, those extra points are eliminated. The literature of building such banks is quite exhaustive, for reference see [89, 90]. The stochastic procedure is described in detail here [91] and it is implemented in `pycbc_aligned_stoch_bank` as part of the PyCBC package [92–94]. Other possible algorithms are `lalapps_cbc_sbank` [95] or novel approaches like `pycbc_brutebank` [96]. Some drawbacks of these methods are that they may take time to converge to the effective coverage required and, on the other side, can tend to over-cover the parameters’ space.

A possible solution is the mixture of these approaches, known as hybrid-methods. Described in [97] and [98], the `hdbank` algorithm proposes a solution that is able to give an efficient hybrid-stochastic placement procedure to generate template banks with aligned spins. The corner stone of this method is to find a coordinate systems where the metric is as flat as possible. This is done by considering a 3-dimensional space in the masses m_1, m_2 and the reduced spin χ_r

$$\chi_r = \left(1 - \frac{76\eta}{113}\right)\chi_s + \delta\chi_a \quad (3.1)$$

Where $\chi_s = \frac{\chi_1^z + \chi_2^z}{2}$ and $\chi_a = \frac{\chi_1^z - \chi_2^z}{2}$ are, respectively, the symmetric and anti-symmetric spins. $\eta = \frac{m_1 m_2}{(m_1 + m_2)}$ is the symmetric mass ratio and $\delta = \frac{m_1 - m_2}{(m_1 + m_2)}$ is the anti-symmetric mass ratio. The algorithm then proceeds as following: in this 3-dimensions space, a uniform grid of few millions points is generated, then a transformation is applied in the dimensionless chirp time coordinates system defined as

$$\theta_0 = \frac{5}{2^{\frac{1}{3}}} \left(\frac{1}{16\pi f_0 (m_1 + m_2) \eta^{\frac{3}{5}}} \right)^{\frac{5}{3}} \quad (3.2)$$

$$\theta_3 = \left(\frac{16\pi^5 \theta_0^2}{25\eta^3} \right)^{\frac{1}{5}} \quad (3.3)$$

$$\theta_{3s} = \frac{113\chi_r \theta_3}{48\pi} \quad (3.4)$$

In this coordinates system, the metric changes more slowly. Then a random point from the grid is randomly chosen and added to the final bank and the local metric is then computed. At this point, a truncated 3-dimensional octahedral lattice is used to estimate the position of new templates according to the estimated metric. Finally, the exact match between these templates and the already selected is computed, and the surviving templates are added to the bank. This procedure is repeated until no more templates are added to the set. This algorithm can produce template banks with 5 – 10% less templates with respect to pure stochastic ones.

The actual MBTA's template banks production consists in generating one bank named 'UBER BANK' that spans the whole parameters' space, from BNS to BBH. For building this, two seeds banks — with generally a higher fitting factor than the UBER BANK – constitute the basis for the UBER BANK. The seeds are created taking into account different procedures for BNS and BBH, the former is created using `pycbc_geom_aligned_bank`, while the latter using `hdbank`. Once these two seeds banks are generated, the final UBER BANK is then created starting from these two, using the `hdbank` algorithm. This procedure is done independently for the full frequency band bank VT and for the two RT banks. In this way, the three different banks are obtained, and then the mapping between those banks must be computed to associate the RTs to the VT, since MBTA will match-filter the data in the Real Template banks, but the trigger is provided by the Virtual Template bank.

The mapping procedure is done by computing the match with the RT templates using injections that are actually the VT bank templates. In this way a map of the matches between the VT and the RTs is done, and the best match correspond to the effective correspondence between the real and the virtual templates. This procedure is not optimal since this involves computing 10 billions matches: in fact a VT usually has in average one million templates for the actual observing run, while the RTs between 10 – 100 thousands. This means that for the whole procedure, creating the seeds, build the UBER BANK, and perform the matches, the time needed is $o(\text{weeks})$. Speed up this process has been a part of this thesis.

3.2 Real Template bank production optimization

We have seen that the association of the RTs with the VT bank comes with a huge computational cost. A possible solution to this problem is to build the RTs directly from the VT, guaranteeing at the same time the same fitting factor.

The procedure starts from the VT bank, which spans the full detector bandwidth. Since the RT banks in MBTA operate over restricted bandwidths (separated in low and high bandwidths by the frequency cut f_{cut} according to Eq. 2.67), we compute the pairwise matches between templates in the VT itself, but restricted to the relevant frequency band of interest. Templates that show a match above a chosen threshold within this band are considered redundant: only one representative template is retained, while others are discarded and their indices are linked to the retained template. In this way the RTs banks and the corresponding VT-RT association map are generated simultaneously in a single run, leading to a faster procedure. In any case, it is not guaranteed that the fitting factor is inherited from the VT bank, and a study in this direction is done in Sec. 3.3. One basic assumption that is taken in this procedure is that the f_{cut} is chosen to ensure that the SNR is approximately equally shared among the low and high frequency bandwidths. In other words, a generic gravitational-wave signal is expected to contribute comparable SNR to both real template banks, so that each band contributes roughly half of the expected SNR for a typical gravitational-wave signal and both frequency bands are given comparable statistical weight in the detection statistic.

3.2.1 Algorithm design and implementation

Let's define the set of N virtual templates VT of the VT bank as $VT = \{VT_i\}_i^N$. As first passage, VT_i is saved as a Real Template and then we compute the match $\langle VT_i|VT_j \rangle_{\Delta f_k}$ in the frequency band Δf_k , with $k = \{LH, HF\}$, $\forall i, j; i \neq j$. Suppose we impose a minimal match of 0.97, then if it happens that

$$\langle VT_i|VT_j \rangle_{\Delta f_k} \geq 0.97 \quad (3.5)$$

VT_j is discarded and its index is saved and associated to VT_i . In this way we record that VT_j corresponds to VT_i in the frequency band Δf_k . Iterating this procedure for each template in the VT bank, for the LF (Δf_{LF}) and the HF (Δf_{HF}) bands we end up with two sets of templates that correspond to the Real Templates banks along with the association file. A scheme of this procedure is depicted in Fig. 3.1.

Computing the matches requires generating the waveforms from the template parameters, which is computationally expensive. Moreover, we expect that no pairs of waveforms will have a match above 0.97 beyond a certain distance in the parameter space, which can be estimated, for instance, through the chirp-mass M_c . To verify this, we considered the BNS parameter space and computed the match between a given template and the subsequent 15,000 templates — ordered by ascending values of chirp-mass — consistent with the standard bank configuration used by MBTA.

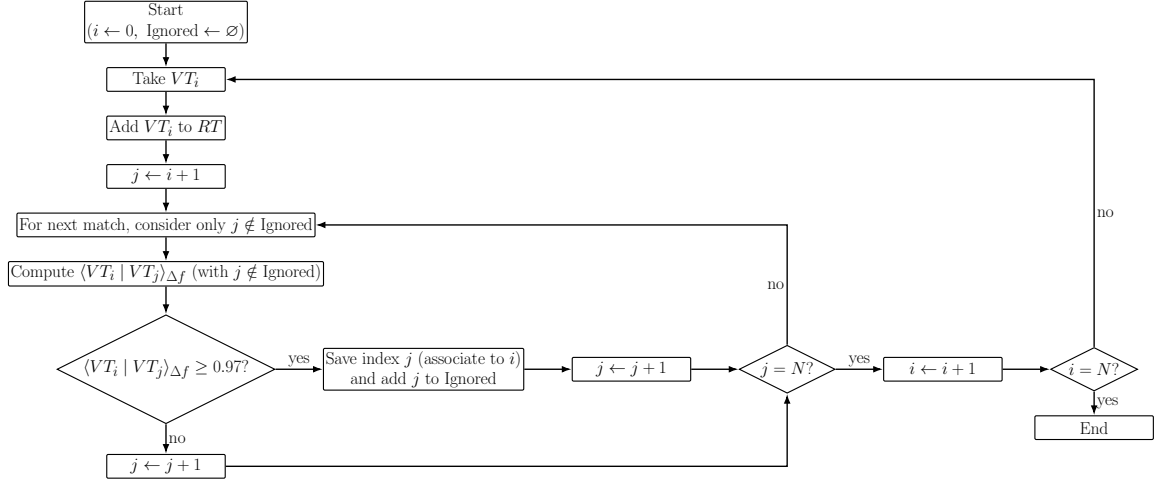



Figure 3.1 Flowchart of the algorithm to construct the RT bank by retaining only non-redundant templates. N is the number of templates. The set Ignored contains indices of redundant templates; these are associated to a real template and excluded from subsequent iterations.

Figure 3.2 shows the distribution of the match as a function of the chirp-mass difference ΔM_c between a reference template and 15 000 neighboring templates in the bank. As expected, the highest matches ($\gtrsim 0.97$) are obtained only for templates with chirp masses very close to that of the reference waveform. This reflects the dominant role of the chirp mass in determining the phase evolution of the inspiral signal.

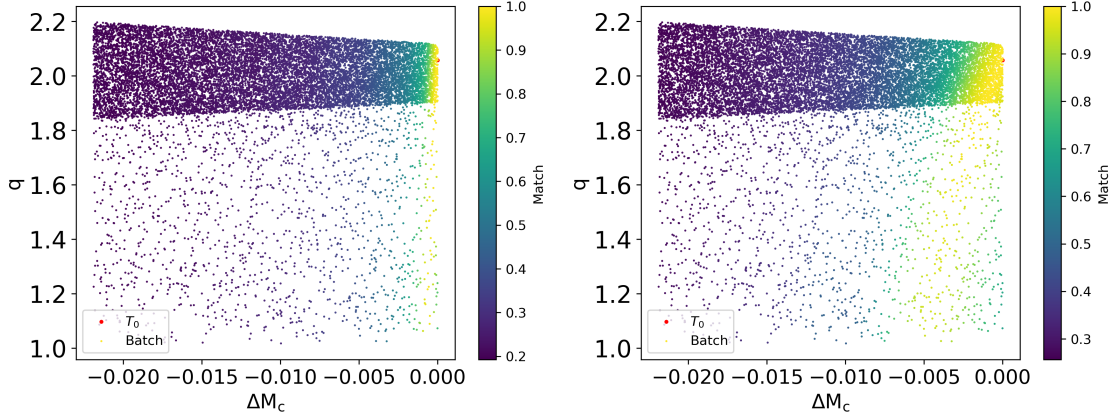
The tails observed as ΔM_c moves away from zero indicate that even small chirp-mass differences rapidly accumulate phase mismatches across the sensitive frequency band, causing the match to drop quickly toward zero. This behavior is expected, since the chirp mass enters the leading-order term of the post-Newtonian phase and therefore strongly constrains the waveform morphology. As a result, templates with slightly different chirp masses become quickly orthogonal to the reference waveform.

Both panels in Fig. 3.2 exhibit a pronounced tail as $\Delta M_c \rightarrow 0$. This behaviour reflects residual mismatches driven by intrinsic parameters other than the chirp mass. In particular, even when the chirp mass is nearly identical to that of the reference template, differences in the mass ratio q can still introduce significant phase and amplitude modulations. These effects prevent the match from saturating uniformly at unity and generate a spread of values close to $\Delta M_c = 0$. We stress that the broader distribution observed in the high-frequency case around $\Delta M_c \simeq 0$ is predominantly driven by variations in the mass ratio q across the template bank. The mass ratio directly affects the detailed morphology of the waveform, with its impact being especially pronounced in the chirp phase, which dominates the high-frequency bandwidth. Consequently, high-frequency matches are more tolerant to chirp-mass degeneracies while remaining sensitive to variations in q , resulting in the wider distribution seen in panel (b).

Consequently, we decided to subdivide of the VT bank — sorted in M_c — into in smaller subsets and compute the matches within these narrower intervals. In addition, this procedure can be naturally



(a) Distribution of matches between the reference template (T_0) and the following 15,000 templates in the low-frequency interval (20–80 Hz). No spins are considered.
 (b) Distribution of matches between the reference template (T_0) and the following 15,000 templates in the high-frequency interval (80–1024 Hz). No spins are considered.



(c) Distribution of matches (colorbar) as function of difference in chirp-mass and mass ratio q .
 (d) Distribution of matches (colorbar) as function of difference in chirp-mass and mass ratio q .

Figure 3.2 Match distributions for a binary neutron star template bank batch as a function of the chirp-mass difference ΔM_c with respect to the reference template T_0 . Panels (a) and (b) show the distributions of matches computed in the low-frequency (20–80 Hz) and high-frequency (80–1024 Hz) bands, respectively, for 15,000 non-spinning templates. Panels (c) and (d) display the match as a function of both ΔM_c and the mass ratio q , with the colour scale indicating the match value. In the low-frequency band, the match is primarily governed by ΔM_c and remains nearly constant across different values of q at fixed chirp mass. In contrast, in the high-frequency band the match exhibits a stronger dependence on q , leading to a broader distribution near $\Delta M_c \simeq 0$. Only a small fraction of templates achieve matches above 0.97, concentrated close to the chirp mass of T_0 .
 Templates are ordered in ascending chirp mass, hence $\Delta M_c < 0$.

parallelized, boosting significantly the computation's speed. The choice of maintaining batches of $\sim 15,000 - 20,000$ templates for the low and high frequency bandwidth is conservative as we expect no templates to exhibit matches above 0.97 at such large separations in parameter space, while this batching strategy ensures efficient use of computational resources during parallel processing.

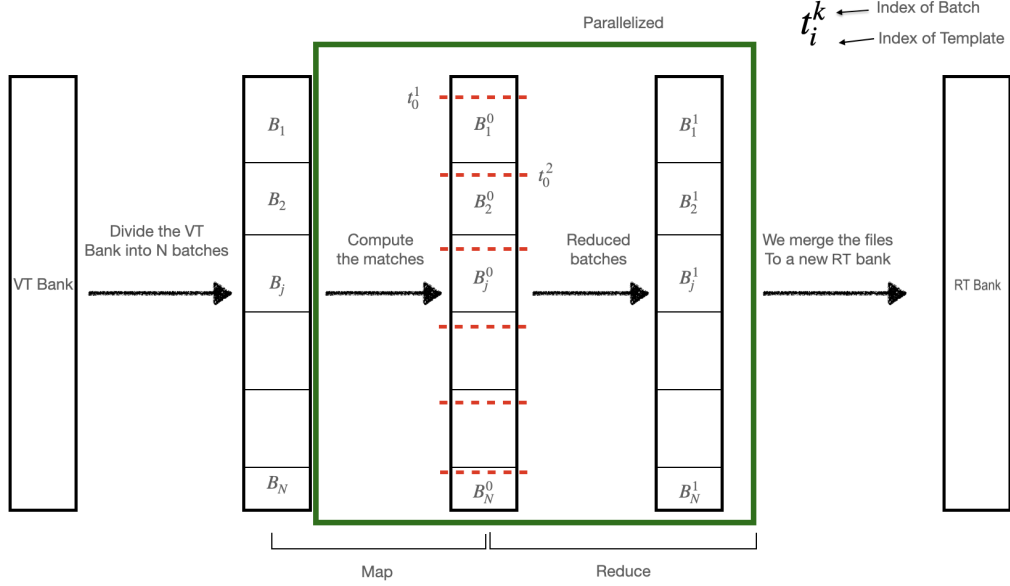


Figure 3.3 Schematic representation of the map–reduce procedure applied to the VT bank to compute the Real Template (RT) banks. Starting from the Virtual Template Bank, we first divide it into N batches (map phase). In each batch, the matches among templates within the considered frequency bandwidth are computed, and when the match exceeds a given threshold, the corresponding template is discarded and its index recorded (reduce phase). Finally, the reduced batches are merged to obtain the final RT bank. For reference, only the first template is highlighted for each batch from which the procedure starts.

To implement this procedure, we build up a code starting from the map–reduce procedure developed in the big-data research area [99]. This approach consists in dividing the original dataset in sub-sets named 'Batches' (map phase), recursively compute the needed operation in these Batches (reduce phase) and finally merge together the results. The scheme of the algorithm is represented in Fig. 3.3.

The codes have been developed using python, in particular the pyCBC [92] libraries for waveforms production and matches computation, while the parallelization procedure is performed by HTCondor at the CIT clusters.

3.2.2 First study: the BNS Virtual Template case

As first study, we considered a BNS virtual-template bank, with parameters usually used by MBTA for building the seed, as reported in Table 3.1 and using the `pycbc_geom_aligned_bank` algorithm. The VT bank template bank obtained in this procedure results in 218,582 templates. The $m_1 - m_2$ plot is reported in Fig. 3.4. The 'hole' in the central value of the masses is an effect induced by the spins that impact the metric computation and results in a sparsity of template in this particular region, even though it does not affect the recovering capabilities of the metric.

BNS-Seed	
Parameter	Value
min-match	0.98
frequency	25–2048 Hz, $\Delta f = 0.01$
m_1	$[1-3] M_\odot$
m_2	$[1-3] M_\odot$
$\min(m_1 + m_2)$	$2 M_\odot$
$\max(m_1 + m_2)$	$4 M_\odot$
χ	$[-0.05, 0.05]$ uniform
ASD	ASDLigo O4a

Table 3.1 BNS production of Virtual Template Bank parameters - ASD reference: [8].

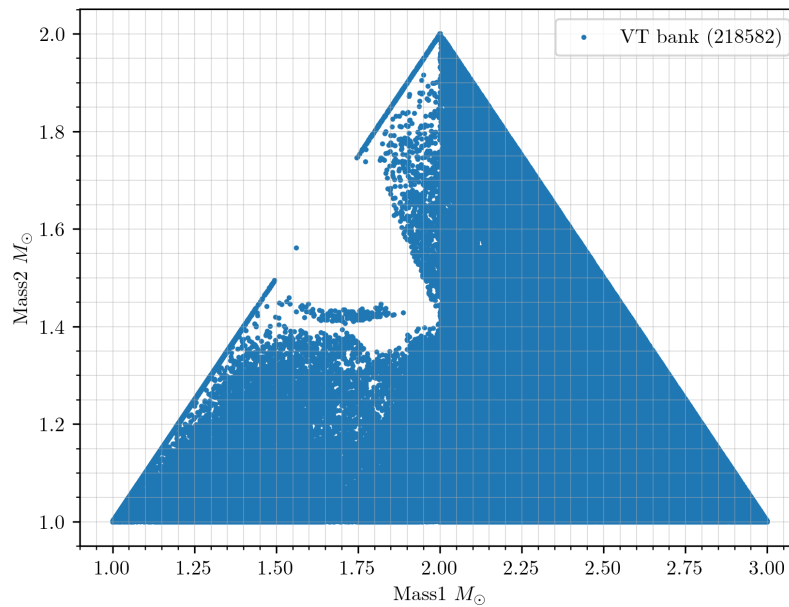


Figure 3.4 Virtual Template bank for the BNS case.

At this point, this VT bank is fed to the map-reduce algorithm described above, where the frequency cut value has to be posed. We recall that this frequency value corresponds to the value where the SNR is equally shared among the bands. To choose the frequency cut we developed a simple algorithm that computes the cumulative SNR^2 of each template as function of frequency. The frequency that corresponds to the 50% of SNR^2 is picked as frequency cut for that given template. An example of such procedure is shown in Fig. 3.6.

Since we are considering templates that are similar in the physical parameter space we expect the frequency cut not to be dramatically different among the different templates. For the whole bank, a single frequency cut is chosen as the median among all the 50% SNR-frequency value and this is

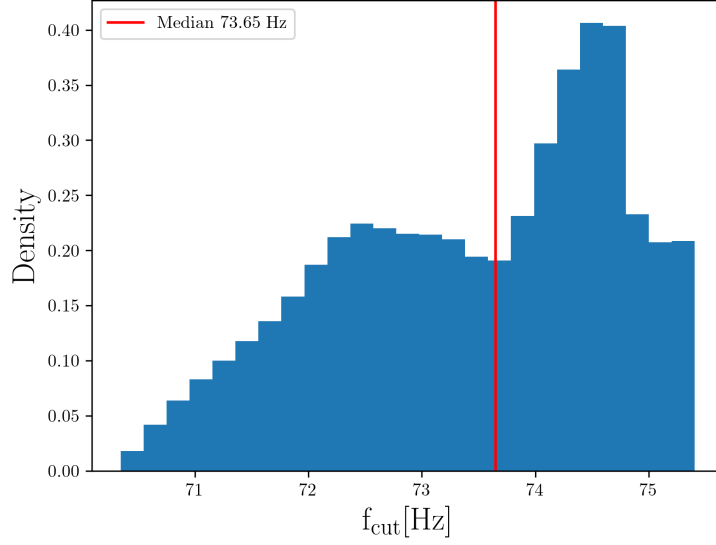


Figure 3.5 f_{cut} distribution for the BNS bank. The median value — highlighted in red — indicated the frequency cut chosen for building the RT banks.

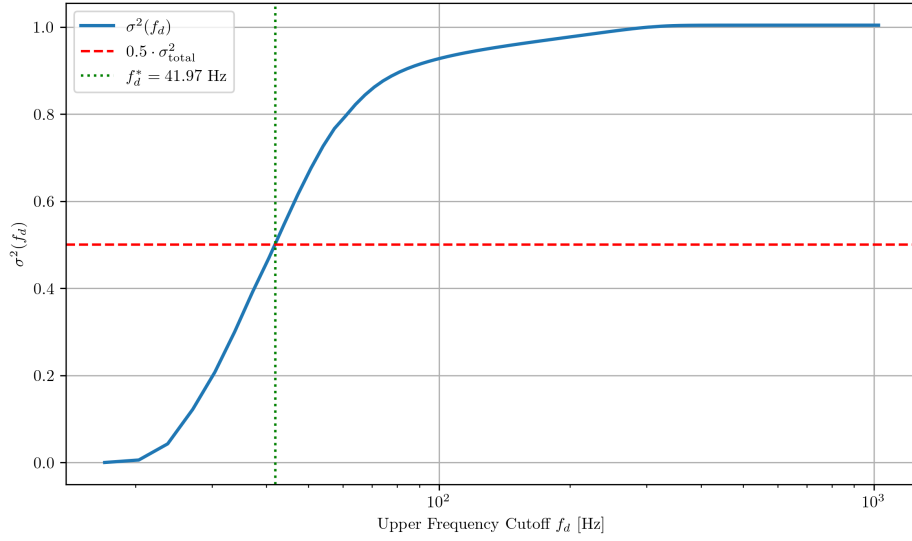


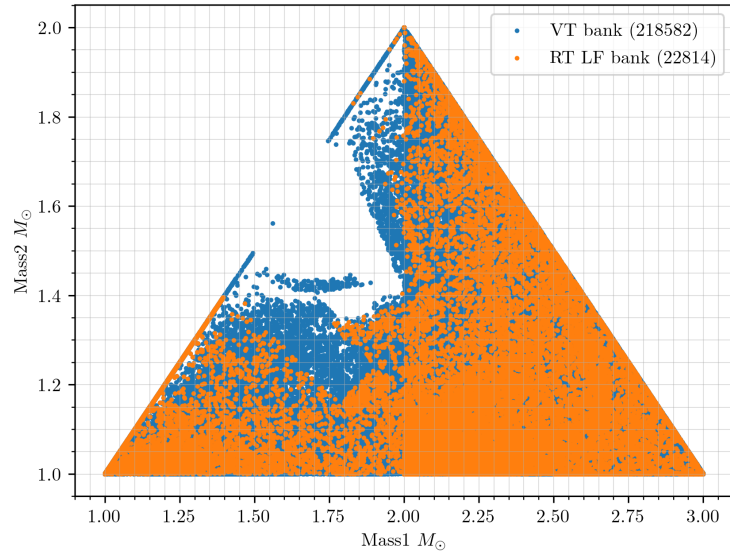
Figure 3.6 Example of procedure for selecting the frequency cut for a single template. In this case a BBH ($m_1 = 46.81M_{\odot} - m_2 = 4.26M_{\odot} - \chi_1^z = \chi_2^z = -0.803$) is taken into account. The blue curve represent the cumulative SNR^2 (σ^2), the red line is the 50% – SNR^2 point and the green-dot line is the corresponding frequency.

$f_{\text{cut}} = 74\text{Hz}$. The corresponding bands are $\text{LF}_{\text{band}} = [25 - 74]\text{Hz}$ – and $\text{HF}_{\text{band}} = [74 - 2048]\text{Hz}$ (see Fig. 3.5).


The noise spectral density considered is the ASDLigo O4a [8], that is the expected sensitivity curve of the LIGO IFOs for the O4a data acquisition campaign, and the approximant used to generate the waveform starting from the templates' parameters is the SEOBNRv4_ROM [100]. The frequency resolution adopted is $\Delta f = \frac{1}{16}$ for the LF band and $\Delta f = \frac{1}{32}$ for the HF band — this difference is due since for high frequencies we need higher sample rate. The minimal match required to eliminate and associate the VT template is posed at 0.98, that is the same value used to construct the VT bank. The batch dimension imposed is 20,000 templates per batch for both bandwidths. The RT banks obtained results in 22,814 templates for the LF bank and 24,581 templates for the HF bank as reported in Fig. 3.7 and the time required is $o(\text{hours})$, along with the association file.

At this point, we can verify that the coverage is guaranteed. The procedure in this sense is done by performing an injection campaign in the parameter space of the bank, compute the best match for each injection among the template bank and verify this value is consistent with the minimal-loss allowed in the bank construction. The matches have been computed using ASDLigo O4a with SEOBNRv4_ROM approximant. For each bank, a total of 1000 injections have been performed. For the RT banks we restricted the analysis to the same frequency bands used in the construction of the bank to compute the matches, while for the VT bank the full bandwidth was used. The results are reported in Fig. 3.8 and Fig. 3.9.

We see that the coverage in the RT banks obtained in this study is guaranteed with no losses below the value of 0.97. We recall that the VT bank was constructed with a minimal match of 0.98, not 0.97, and approximately 10% of the injections lie in this 0.97-0.98 interval (see Fig. 3.8b). This means that although the nominal fitting factor is 0.98, the procedure used to construct the VT bank did not fully converge to that value, and the RT banks inherit this property. In any case, we observe that only a residual fraction on injections falls below 0.98. This result suggests that the RT banks obtained in this way inherit the coverage from the original VT bank. In the next section we propose a more formal explanation of this property.

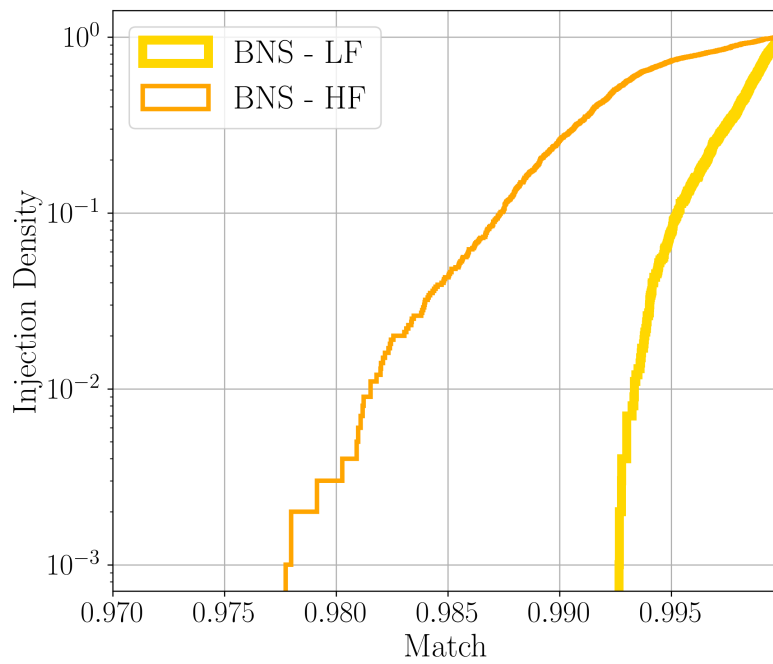


(a) RTL bank

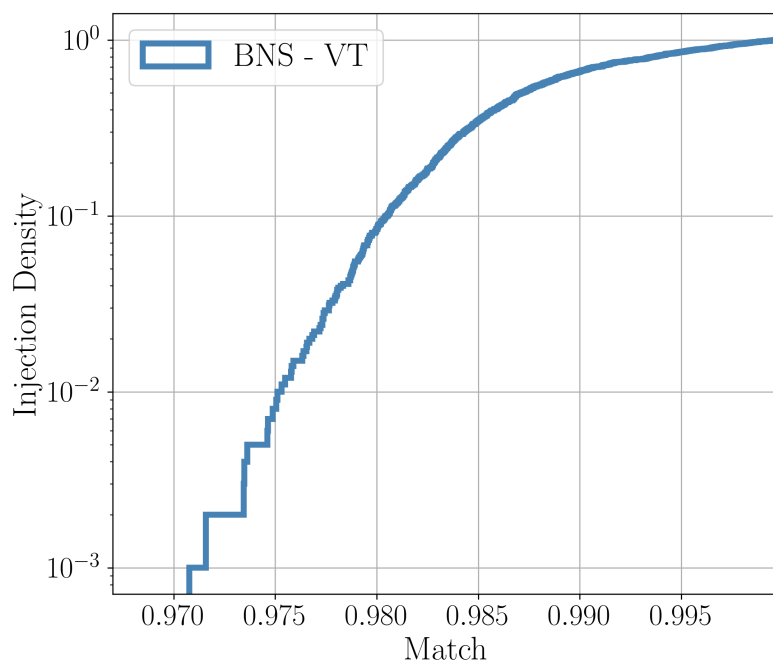


(b) RTHF bank

Figure 3.7 Map-reduce procedure applied to the VT bank for computing the RT banks result.




(a) Cumulative distribution for the RT banks fitting factor (match).

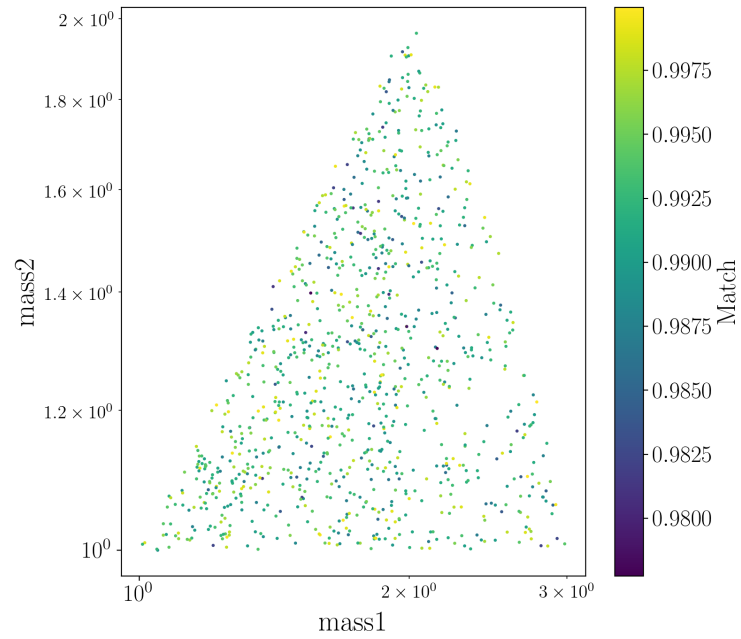


(b) Cumulative distribution VT bank fitting factor (match).

Figure 3.8 Bank Vericator procedure applied to the RTLF and RTHF banks (a) and the BNS VT bank (b)



(a) Scatter plot - RTL match



(b) Scatter plot - RTHF match

Figure 3.9 Scatter plot for injected masses. The colorbar represents the match gradient for each injection. We see that no particular behavior emerges from the distribution, the injections are uniformly recovered within the expected fitting factor.

3.3 The analytical coverage

In this section we want to provide a solid argument for the procedure described above. We want to show that given a bank which guarantees us a coverage of any signal $h(t)$ — let's say in this case 0.97 — in whole frequency bands, the real template banks obtained from the same bank, considering the cut in frequency, have a coverage compatible with this, with an inferior limit determined by the triangular inequality. For completeness, we recall some definitions. We define the inner product between two function of time $a(t)$ and $b(t)$ as

$$\langle a|b \rangle = 4Re \int_0^\infty df \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} \quad (3.6)$$

Where $\tilde{a}(f)$ is the Fourier Transform of $a(t)$ defined as

$$\tilde{a}(f) = \int_{-\infty}^\infty dt e^{2i\pi ft} a(t) \quad (3.7)$$

This scalar product can also be defined as the match between a and b . Actually, in the Gravitational Waves study field, also the extrinsic variables as phase ϕ , amplitude A and time t should be taken into account in the evaluation, but those can be integrated out [21], so we can proceed with the current notation. Clearly if we consider the scalar product between two functions in a specific frequency range Δf , we have

$$\langle a|b \rangle_{\Delta f} = 4Re \int_{\Delta f} df \frac{\tilde{a}^*(f)\tilde{b}(f)}{S_n(f)} \quad (3.8)$$

Where Δf is the interval $(f_1, f_2); f_1 < f_2$. In general, a scalar product induces a geometry in the space, geometry that is defined by a metric. Let's consider for instance now a function $\tilde{u} = \tilde{u}(f; \bar{\lambda})$, where $\bar{\lambda}$ is a vector of parameters which defines the function \tilde{u} . If we consider now the match between two functions $\tilde{u}(f; \bar{\lambda})$ and $\tilde{u}(f; \bar{\lambda} + \Delta\bar{\lambda})$ we can introduce the following function

$$M(\bar{\lambda}, \Delta\bar{\lambda}) = \langle u(\bar{\lambda})|u(\bar{\lambda} + \Delta\bar{\lambda}) \rangle_{\Delta f} \quad (3.9)$$

This quantity, that is the match, is also known as the ambiguity function [71], and it tells us how similar two functions are in the parameters' space defined by $\bar{\lambda}$. Clearly the match has the maximum value around $\Delta\bar{\lambda} = 0$, so we can expand this function around this point obtaining,

$$M(\bar{\lambda}, \Delta\bar{\lambda}) \simeq 1 + \frac{1}{2} \left(\frac{\partial^2 M}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right) \Big|_{\Delta\bar{\lambda}=0} \Delta\lambda^i \Delta\lambda^j \quad (3.10)$$

So we can define the metric

$$g_{i,j} = -\frac{1}{2} \left(\frac{\partial^2 M}{\partial \Delta\lambda^i \partial \Delta\lambda^j} \right) \Big|_{\Delta\bar{\lambda}=0} \quad (3.11)$$

Then the mismatch between two functions is

$$1 - M = g_{i,j} \Delta \lambda^i \Delta \lambda^j \quad (3.12)$$

which is basically a distance. So given two functions h_1 and h_2 , the distance d_{h_1, h_2} between these functions is

$$d_{h_1, h_2} = 1 - \langle h_1 | h_2 \rangle_{\Delta f} \quad (3.13)$$

Another fundamental quantity is the normalization of a function, which means that if we want that a function h in the frequency interval Δf is normalized to one

$$\langle h | h \rangle_{\Delta f} = 1 \quad (3.14)$$

We need to impose that

$$h = \frac{\hat{h}}{\sqrt{\langle \hat{h} | \hat{h} \rangle_{\Delta f}}} \quad (3.15)$$

where \hat{h} is the unnormalized function. We must specify that the normalization is performed among the frequency band defined in the scalar product.

3.3.1 Relationship between the match and the frequency bands

Let's consider now two normalized waveforms h_1 and h_2 and compute the scalar product:

$$\langle h_1 | h_2 \rangle_{\Delta f} = 4Re \int_{\Delta f} df \frac{\tilde{h}_1^*(f) \tilde{h}_2(f)}{S_h(f)} \quad (3.16)$$

we can explicit the expression of such normalized functions as

$$h_1 = \frac{\hat{h}_1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}}} \quad (3.17)$$

$$h_2 = \frac{\hat{h}_2}{\sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \quad (3.18)$$

consider now the following frequency intervals

$$\Delta f = (f_1, f_2)$$

$$\Delta f_1 = (f_1, f_c)$$

$$\Delta f_2 = (f_c, f_2)$$

with $f_1 \leq f_c \leq f_2$. The scalar product between h_1 and h_2 can be expressed as sum of the different frequency interval, since the linear property of integrals, and explicitly writing the normalized functions results in

$$\begin{aligned}
\langle h_1 | h_2 \rangle_{\Delta f} &= 4\text{Re} \int_{\Delta f_1} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} + 4\text{Re} \int_{\Delta f_2} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} = \\
&= \frac{1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} 4\text{Re} \int_{\Delta f_1} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} + \frac{1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} 4\text{Re} \int_{\Delta f_2} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} = \\
&= \frac{1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_1} + \frac{1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_2}
\end{aligned} \tag{3.19}$$

Focusing on the normalized match in the corresponding frequency intervals — considering for instance Δf_1 — we see that

$$\langle h_1 | h_2 \rangle_{\Delta f_1} = 4\text{Re} \int_{\Delta f_1} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} = \frac{1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}} 4\text{Re} \int_{\Delta f_1} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} \tag{3.20}$$

clearly

$$\langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_1} = 4\text{Re} \int_{\Delta f_1} df \frac{\tilde{h}_1^* \tilde{h}_2}{S_h(f)} \tag{3.21}$$

then

$$\langle h_1 | h_2 \rangle_{\Delta f_1} = \frac{1}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}} \langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_1} \tag{3.22}$$

Inverting this expression for $\langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_1}$ and putting all together in Eq. 3.19 we have

$$\langle h_1 | h_2 \rangle_{\Delta f} = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \langle h_1 | h_2 \rangle_{\Delta f_1} + \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_2}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \langle h_1 | h_2 \rangle_{\Delta f_2} \tag{3.23}$$

In order to slim the notation we define

$$\alpha = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \quad (3.24)$$

$$\beta = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_2}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \quad (3.25)$$

then we can finally write

$$\langle h_1 | h_2 \rangle_{\Delta f} = \alpha \langle h_1 | h_2 \rangle_{\Delta f_1} + \beta \langle h_1 | h_2 \rangle_{\Delta f_2} \quad (3.26)$$

3.3.2 Properties of α and β

Let's start considering α . We know that $\langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_1} \leq \langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f}$, since

$$4\text{Re} \int_{\Delta f_1} df \frac{\tilde{h}_1^* \tilde{h}_1}{S_h(f)} \leq 4\text{Re} \int_{\Delta f} df \frac{\tilde{h}_1^* \tilde{h}_1}{S_h(f)} \quad (3.27)$$

If \tilde{h}_1^* and \tilde{h}_1 are functions defined positive, it follows that the integrals above are quantities equal or greater than zero, then 3.27 holds. This condition is acquainted since $\tilde{h}_1^* \tilde{h}_1 = |\tilde{h}_1|^2$, which is positive. Similarly we have $\langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f_2} \leq \langle \hat{h}_1 | \hat{h}_2 \rangle_{\Delta f}$. It follows then that

$$\alpha = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} \leq \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} = 1 \quad (3.28)$$

Then

$$\alpha \leq 1 \quad (3.29)$$

$$\beta \leq 1 \quad (3.30)$$

Now, by interval addition properties of integrals is trivial that

$$\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} = \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} + \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2} \quad (3.31)$$

Let's now suppose that the match between the waveform \hat{h} and itself is equally shared among the frequency bands,

$$\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} = \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2} \quad (3.32a)$$

$$\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1} = \langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_2} \quad (3.32b)$$

This means that if we consider for instance the \hat{h}_1 function,

$$\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f} = \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} + \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2} = 2 \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1} = 2 \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2} \quad (3.33)$$

The result is symmetric for \hat{h}_2 . In this way α can be written as

$$\alpha = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}}{\sqrt{2 \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_1}} \sqrt{2 \langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_1}}} = \frac{1}{2} \quad (3.34)$$

The same procedure can be applied to β

$$\beta = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_2}}}{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f}}} = \frac{\sqrt{\langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2}} \sqrt{\langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_2}}}{\sqrt{2 \langle \hat{h}_1 | \hat{h}_1 \rangle_{\Delta f_2}} \sqrt{2 \langle \hat{h}_2 | \hat{h}_2 \rangle_{\Delta f_2}}} = \frac{1}{2} \quad (3.35)$$

3.3.3 Match sharing among the frequency bands

We imposed that the match is the same between \hat{h} and itself among the frequency bands Δf_1 and Δf_2 , see Eq. 3.32a and 3.32b— the relation is true also for the normalized functions, just insert 3.18 in the relations Eq. 3.32 — it means that under this assumption

$$\langle h_1 | h_1 \rangle_{\Delta f} = \langle h_1 | h_1 \rangle_{\Delta f_1} + \langle h_1 | h_1 \rangle_{\Delta f_2} = 2 \langle h_1 | h_1 \rangle_{\Delta f_1} \quad (3.36)$$

We also showed that Eq. 3.32 imply that $\alpha = \beta$, so Eq. 3.26 is

$$\langle h_1 | h_2 \rangle_{\Delta f} = \alpha \langle h_1 | h_2 \rangle_{\Delta f_1} + \beta \langle h_1 | h_2 \rangle_{\Delta f_2} = \alpha (\langle h_1 | h_2 \rangle_{\Delta f_1} + \langle h_1 | h_2 \rangle_{\Delta f_2}) \quad (3.37)$$

Now, we suppose that h_1 and h_2 are 'similar', at least in Δf , so that their scalar product is almost equal to one

$$\langle h_1 | h_2 \rangle_{\Delta f} \sim 1 \quad (3.38)$$

So Eq. 3.37 became

$$1 = \alpha (\langle h_1 | h_2 \rangle_{\Delta f_1} + \langle h_1 | h_2 \rangle_{\Delta f_2}) \quad (3.39)$$

Including $\alpha = \beta = \frac{1}{2}$ Eq. 3.26 finally results in

$$1 = \frac{1}{2}(\langle h_1|h_2 \rangle_{\Delta f_1} + \langle h_1|h_2 \rangle_{\Delta f_2}) \quad (3.40)$$

Recalling that h is normalized, then this expression can be true if and only if

$$\langle h_1|h_2 \rangle_{\Delta f_1} = \langle h_1|h_2 \rangle_{\Delta f_2} = 1 = \langle h_1|h_2 \rangle_{\Delta f} \quad (3.41)$$

This result has an important consequence: we obtained that the match between the functions among the whole frequency band — if Eq. 3.38 is respected — and the sub-bands can be the same, as the previous equation states.

3.3.4 Coverage problem among different frequency bands

Let us consider a generic Gravitational Wave's signal h and a set of functions (the Virtual Templates) $\mathbb{T} = \{T_k\}_{k=1}^N$ such that for any physically admissible signal h there exists at least one template $T_k \in \mathbb{T}$ satisfying

$$\langle h|T_k \rangle_{\Delta f} \geq 0.97 \quad (3.42)$$

In this situation we say that the coverage provided by the set \mathbb{T} in the frequency band Δf is 0.97, in symbols $\mathcal{C}_{\Delta f}^{\mathbb{T}} = 0.97$

Now, let consider frequency $f_c | f_1 < f_c < f_2$ which defines the following frequency intervals $\Delta f = (f_1, f_2), \Delta f_1 = (f_1, f_c), \Delta f_2 = (f_c, f_2)$. Consider now two templates T_l and $T_k \in \mathbb{T}$ and evaluate their band-limited match

$$\langle T_k|T_l \rangle_{\Delta f_1} \quad (3.43)$$

If it happens that $\langle T_k|T_l \rangle_{\Delta f_1} \geq 0.97$ then the template T_l is considered redundant with respect to T_k in Δf_1 . Accordingly, this non-redundant template T_k is stored into a new set $\bar{\mathbb{T}}$

$$\bar{\mathbb{T}} = \{T_i\}_i^M, M \leq N \quad (3.44)$$

which collects the representatives of the Real Template bank. This pruning procedure is repeated iteratively until no remaining pair of templates in \mathbb{T} satisfies

$$\langle T_i|T_j \rangle_{\Delta f_1} > 0.97, \forall T_i, T_j \in \mathbb{T} \quad (3.45)$$

applying this procedure iteratively, the resulting subset of templates $\bar{\mathbb{T}} = \{T_i\}_i^M \in \mathbb{T}$ is such that all the templates satisfy the condition

$$\langle T_l|T_m \rangle_{\Delta f_1} \leq 0.97 \quad \forall T_l, T_m \in \bar{\mathbb{T}}. \quad (3.46)$$

The resulting relation $\bar{T} \in T$ expresses that the final RT bank is a subset of the original VT bank. For clarity, we report an example of such set: let $T = \{T_1, T_2, T_3, T_4, T_5\}$, then we evaluate the pairwise matches $\langle T_i | T_j \rangle$ within the frequency band Δf_1 . If it happens for this pair that the match exceeds a given threshold, then the j template is considered redundant and associated to the i template. The procedure is iterated for each template in the set.

$$\begin{aligned} \langle T_1 | T_2 \rangle_{\Delta f_1} &\leq 0.97 \\ \langle T_1 | T_3 \rangle_{\Delta f_1} &\geq 0.97 \rightarrow \text{pop}(T_3) \\ \langle T_1 | T_4 \rangle_{\Delta f_1} &\leq 0.97 \\ \langle T_1 | T_5 \rangle_{\Delta f_1} &\geq 0.97 \rightarrow \text{pop}(T_5) \\ \langle T_2 | T_4 \rangle_{\Delta f_1} &\geq 0.97 \rightarrow \text{pop}(T_4) \end{aligned}$$

where $\text{pop}(\cdot)$ means remove T_j from \bar{T} . After this pruning procedure we obtain

$$\begin{aligned} T &= \{T_1, T_2, T_3, T_4, T_5\} \\ \bar{T} &= \{T_1, T_2\} \end{aligned}$$

Eq. 3.46 hold by construction. Then we ask ourselves which relationships exist between h and the functions in the reduced set of templates $\bar{T} = \{T_i\}_i^M$. In particular, does the reduced set preserve the same coverage in Δf_1 ? In symbols:

$$\mathcal{C}_{\Delta f_1}^{\bar{T}} = 0.97 \quad (3.47)$$

Formally, this correspond to verifying whether $\forall h, \exists T_l \in \bar{T} | \langle h | T_l \rangle_{\Delta f} \geq 0.97$ is true. Since we know that $\bar{T} \subseteq T$ the first trivial case is that $T_l = T_k$ for some k (in our example we can put $T_k = T_l = T_1$), in this case we have necessary that

$$0.97 \leq \langle h | T_k \rangle_{\Delta f} = \langle h | T_l \rangle_{\Delta f_1} \quad (3.48)$$

This is true since Eq. 3.42 and 3.41. But what happens when $T_l \neq T_k$ (for instance, $T_k = T_3, T_l = T_1$)? In this case we know that $\forall h \exists T_k \in T$ such that

$$\langle h | T_k \rangle_{\Delta f} = \langle h | T_k \rangle_{\Delta f_1} \geq 0.97 \quad (3.49)$$

Moreover, by construction $\forall T_k \in T$ exists a representative $T_l \in \bar{T}$ such that

$$\langle T_l | T_k \rangle_{\Delta f_1} \geq 0.97 \quad (3.50)$$

At this point we can recall that we are working in a metric space, so we can apply the triangular inequality. In general, given three elements T_1, T_2, T_3 of the metric space, the triangular inequality states that

$$d_{T_1, T_3} \leq d_{T_1, T_2} + d_{T_2, T_3} \quad (3.51)$$

So by Eq. 3.49 and 3.50 we can write the distances as

$$d_{h, T_k} = 1 - \langle h | T_k \rangle_{\Delta f_1} \leq 1 - 0.97 = 0.03 \quad (3.52)$$

$$d_{T_l, T_k} = 1 - \langle T_l | T_k \rangle_{\Delta f_1} \leq 1 - 0.97 = 0.03 \quad (3.53)$$

resulting in

$$d_{h, T_l} \leq d_{h, T_k} + d_{T_k, T_l} \leq 0.03 + 0.03 = 0.06 \quad (3.54)$$

and this is true $\forall, T_k \in \mathbb{T}$ and their corresponding $T_l \in \bar{\mathbb{T}}$. This means that a Coverage $\mathcal{C}_{\Delta f_1}^{\bar{\mathbb{T}}} = 0.94$ is guaranteed for this kind of bank. This result shows that, under the assumption that the total match (or equivalently the signal's SNR) is equally shared between the two frequency bands, the Real Template banks built using this procedure retain a guaranteed coverage of at least 0.94 — a direct consequence of the triangle inequality in the metric space of waveform overlaps. However, different templates may have different f_{cut} values, so the equal-SNR assumption does not hold in practice. To address this — as we have seen in the BNS case — we approximate the split-frequency by using the median f_{cut} over the distribution of frequencies computed in the VT bank.

3.4 The UBER BANK case

In this section we study the UBER BANK case, that is the template bank usually used by MBTA for a search that involves BNS, BBH and NSBH. The procedure to build such bank, as described at the beginning of this chapter, consists in constructing at first two smaller banks known as 'Seed banks', one for the BNS and one for the BBH, considering only a fraction of the parameter space but imposing a fitting factor usually bigger than the one required for the whole bank. This is done since during O3 the great majority of events fell in this region of the parameter space, so we tighten-up the minimal match required to be sure not to lose SNR. The BNS Seed is obtained via `pycbc_geom_aligned_bank` algorithm, while the BBH Seed is from the `hdbank` procedure. Then the Seed banks are given to the `hdbank` algorithm for the final UBER BANK. The banks parameters are reported in Table 3.2.

BNS Seed	
Parameter	Value
min-match	0.98
f	25–2048 Hz, $\Delta f = 0.01$
m_1	[1–3] M_\odot
m_2	[1–3] M_\odot
$\min(m_1 + m_2)$	2 M_\odot
$\max(m_1 + m_2)$	4 M_\odot
χ	[-0.05, 0.05] uniform
ASD	ASD-O4-Ligo

(a) BNS-Seed production Virtual Template Bank parameters

BBH Seed	
Parameter	Value
min-match	0.98
f	18–1024 Hz, $\Delta f = 0.1$
m_1	[5–100] M_\odot
m_2	[5–100] M_\odot
$\min(m_1 + m_2)$	10 M_\odot
$\max(m_1 + m_2)$	199 M_\odot
χ	[-0.998, 0.998] uniform
ASD	ASD-O4-Ligo

(b) BBH-Seed production Virtual Template Bank parameters

UBER-BANK	
Parameter	Value
min-match	0.965
f	18–1024 Hz, $\Delta f = 0.1$
m_1	[1–100] M_\odot
m_2	[1–100] M_\odot
$\min(m_1 + m_2)$	2 M_\odot
$\max(m_1 + m_2)$	199 M_\odot
χ_{BNS}	[-0.05, 0.05] uniform
χ_{BBH}	[-0.998, 0.998] uniform
ASD	ASD-O4-Ligo

(c) UBER BANK Template Bank parameters

Table 3.2 Comparison of Virtual Template Bank Parameters

The final UBER BANK results then in 896,634 templates. To construct the RT banks, we look for the median frequency cut among the full UBER BANK and it results then in $f_{\text{cut}} = 70\text{Hz}$, see Fig. 3.10 for reference. We also took into account the different minimal match required for the BBH and BNS Seeds. Accordingly, we imposed as threshold to eliminate the redundant templates respectively 0.98 for the seeds and 0.965 for the rest of the bank. This procedure resulted finally in 67,404 templates for the LF band and 59,849 templates for the HF (see Fig. 3.11). The computational time required to produce the RT banks along with the association file is $o(\text{hours})$.

We now verify the fitting factor of the banks, with analogous injection campaign as done for the BNS-only case but considering this time also the BBH and NSBH parameter space with 1000 injections each in the corresponding bandwidths for the VT and the RTs. Results of the cumulative fitting-factor distributions are reported in Fig. 3.12 and 3.13.




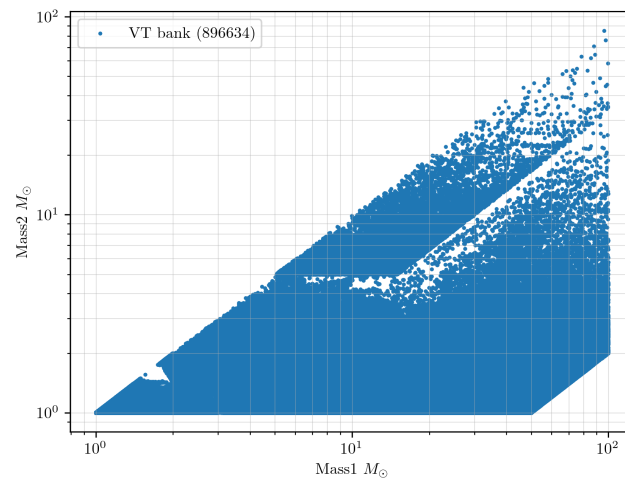
Figure 3.10 Distribution of f_{cut} among the UBER BANK. The red line identifies the median frequency cut and correspond to the chosen frequency for bandwidth selection.

The fitting factor provided by the RTLF is consistent with expectations. Only a small fraction of injections—about $\sim 1\%$ of high-mass BBHs—are recovered with a match below 0.97. These outliers are primarily associated with boundary effects in the hdbank placement during the VT bank construction. Such effects tend to impact the coverage near the upper edge of the mass range — see Fig. 3.14 for reference. Conversely, the RTHF exhibits good coverage across the entire parameter space.


It is important to remark that, in any case, the number of templates in the RT banks produced with this procedure is significantly larger than the standard procedure. In fact, the classical Real Template bank generation for this region of the parameter space typically results in a few tens of thousands of templates, usually around 20,000 for the HF and 30,000 for the LF. In contrast, the present method yields approximately twice as many: 67,404 for the LF and 59,849 for the HF. The main advantage of this approach lies in the speed of generating both the banks and the association files. Nevertheless, further optimization is required to reduce the number of templates, as the current size has a significant impact on computational cost of matched-filtering.

3.4.1 Impact of the f_{cut} choice


In Sec. 3.2.2 we constructed a real template bank starting from a virtual one by computing a frequency cut that defines two bands in which the SNR is equally shared. In the BNS systems where the frequency cut is similar for all the templates the median value approximation for the f_{cut} is quite good. The problem becomes more complex if we instead take into account the BBH or NSBH systems. For more massive or asymmetric systems the signals will be much shorter, and the contribution to



(a) UBER BANK




(b) RTLF bank




(c) RTHF bank


Figure 3.11 Map-Reduce algorithm results for the VT UBER BANK to obtain the RT Banks.



(a) UBER BANK - RTLF




(b) UBER BANK - RTHF




(c) UBER BANK - full bandwidth

Figure 3.12 Cumulative fitting factor obtained for the RT banks and the UBER BANK in the whole parameter space of a full search.



(a) Scatter plot - RTL match



(b) Scatter plot - RTH match

Figure 3.13 Scatter plot for injected masses. The colorbar represents the match gradient for each injection. Even though the coverage is within the expected values, we see that the degradation of the fitting factor associated to the high-mass parameter space of BBH in the LF band.

the SNR is given for a consistent fraction from the chirp phase of the binary evolution. This means that the frequency cut can be significantly lower than the BNS case, as Fig. 3.6 —where a BBH is studied— shows that $f_{\text{cut}} \sim 40\text{Hz}$. The choice of 70Hz for all the templates, as done in this case, can lead to a non-equal distribution of SNR among the bands, consequently the method just shown may result in non-covering real template bank. Investigating in this direction, we found out that that the symmetric mass ratio q deeply impact the f_{cut} selection . In Fig. 3.15 is depicted this kind of behavior.




Figure 3.14 Scatter plot of injected masses with Fitting Factor (Match) for the full-bandwidth Virtual Template bank (Focus on BBH parameter space). We see that at top right corner the VT bank fails to recover properly the injected signals, this means that the RT banks that are constructed from this one will inherit the same coverage behavior.

This parameter affects the waveform morphology, modifying consequently the frequency value where the SNR is equally shared, and in the following we will show this effect through an example.

In general, considering the UBER BANK case studied, this bank yields an average mass ratio $\hat{q}^{\text{UB}} = 9.8$ and an average cut frequency $\hat{f}_{\text{cut}}^{\text{UB}} = 70\text{Hz}$, where the $\hat{\cdot}$ symbol identifies an average quantity. If we instead consider a representative mass ratio $q = 17$, the corresponding cut frequency decreases to $f_{\text{cut}} = 43\text{Hz}$. These differences can have a significant impact on the signal-to-noise-ratio distribution across the frequency band.

As a case study we considered a template T_0 where these effects are particularly evident. This template T_0 has $m_1 = 74.1M_{\odot}$, $m_2 = 4.33M_{\odot}$, and $q = 17$; for simplicity, spin effects are neglected.

When computing the SNR obtained by matched-filtering the signal with itself in the respective frequency intervals, we find that for a cut frequency of $f_{\text{cut}} = 70\text{Hz}$, the SNR in the range $20 - 70\text{Hz}$ is 105, while in the interval $70 - 1024\text{Hz}$ we obtain 61. That implies that roughly two third of the signal power lies in the LF band, and the assumption that the SNR must be equally shared among the bands is no longer valid, challenging the procedure used to built the RT banks.





Figure 3.15 Distribution of f_{cut} among the UBER BANK in the $q - \chi_{\text{eff}}$ space. The colorbar shows a gradient with descending frequency values toward high q and low χ_{eff} values.

Conversely, if we adopt as frequency-cut the one associated with the template itself, $f_{\text{cut}} = 43\text{Hz}$ we obtain that an SNR of 85 in both LF and HF bands. For reference, see Fig. 3.16. The waveforms are build using the SEOBNRv4_ROM approximant in the frequency interval specified for this SNR study.


To corroborate the hypothesis that the observed imbalance in SNR distribution is primarily driven by the high mass-ratio, we considered a comparison signal with same chirp-mass of T_0 and computed the SNR distribution between the two frequency bands with $f_{\text{cut}} = 70\text{Hz}$. The chirp-mass of T_0 is $M_{\text{c},T_0} = 13.2$. We then define a symmetric system T_1 with $m_1 = m_2 = 15.2M_{\odot}$, yielding $M_{\text{c},T_1} = 13.2$. Spins were neglected also in this case.

The resulting SNR in the 20 – 70Hz and 70 – 1024Hz bands for the T_1 case are of 85 and 86, respectively, confirming that for systems with symmetric masses the SNR is nearly evenly distributed across the bands, and the imbalance observed for T_0 is indeed associated with its large mass ratio.

Although in this study we found that the imbalance in the SNR distribution does not appear to have a severe impact on the injection recovery capabilities of the Real Template banks, the choice of the frequency cut remains an important parameter that must be carefully tuned. In particular, we observe that the nominal value $f_{\text{cut}} = 80\text{Hz}$, adopted by MBTA is not optimal in terms of equally splitting the SNR among the frequency bands — as reported in Fig. 3.17 — and can lead to suboptimal matched-filtering performance. In such cases, the signal-to-noise ratio of a given signal may be



(a)



(b)

Figure 3.16 Pictorial representation of the impact of q for T_0 on the SNR distribution among the frequency bands. On top (a) the $f_{\text{cut}} = 70\text{Hz}$ shows how the SNR distribution is unbalanced in favor of the LF band while in (b) the frequency cut imposed is $f_{\text{cut}} = 43\text{Hz}$ with an SNR equally shared among the bands. The waveforms are obtained using SEOBNRv4_ROM in the frequency interval specified for the SNR study.

predominantly confined to one frequency band, potentially preventing a coincident trigger between the low- and high-frequency templates.

Regarding the Real Template bank production, we observe that even if the time required in producing these is substantially reduced, significant increase of template following this method with respect to the state-of-the-art procedure is affecting the procedure. One hypothesis for this behavior is that the selected templates from the VT bank are not optimally placed if we consider the respective frequency interval for the low and high frequency. It is known that the metric depends from the bandwidth we are considering, consequently affecting the template placement [97]. Taking into

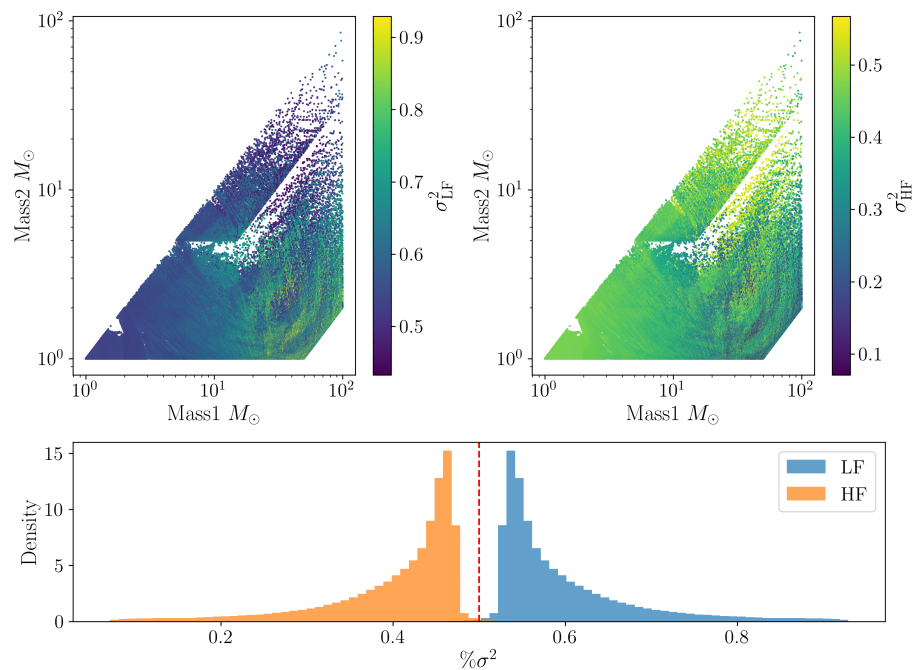


Figure 3.17 Top: scatter plot with UBER BANK masses and corresponding percentage of SNR percentage for each template (LF - left; HF - right). Bottom: distribution of SNR percentage for the two frequency bands. The analysis is performed using the nominal frequency cut of 80 Hz adopted by MBTA.

account this can significantly improve the method, for instance by selecting, or even posing, templates in points of the space that ensure a reduction of number of template with respect to the presented method.

Chapter 4

Random Forests for MBTA

Gravitational-wave science is a highly complex and interdisciplinary field, involving challenges such as detector noise mitigation, source modeling, signal detection and parameter estimation. Comprehensive reviews of machine learning (ML) applications in this domain can be found in [101, 102] highlighting how ML has become an increasingly widespread and powerful tool in GW astronomy.

As outlined in Sec. 2, multiple stages are required to process interferometric data and infer the properties of astrophysical sources. ML techniques have been integrated into each of these stages to improve efficiency and performance. For example, significant progresses have been made in the area of noise mitigation. Classical data analysis techniques assume stationary, Gaussian noise, whereas real detector data often deviate from this premise due to non-stationary behavior and transient phenomena. The term 'non-stationarity' refers to slow shifts in the noise floor, while 'glitches' describes short-duration bursts of excess noise. Detecting, characterizing and classifying such noise artifacts benefited a lot from machine learning techniques [103–109].

It is worth noting the substantial body of work devoted to the noise characterization in the detector data strain. In this context, ML algorithms have made significant contributions in classification and prediction modeling of detector noise in time domain, as well as to the subtraction of noise lines in spectral representations. One notable example is the Gravity Spy algorithm [110–112], which provides an effective tool for classifying glitches present in the data strain. This method, based on a neural-network architecture, takes as input the Omega scans of glitches and categorizes them into several classes corresponding to known glitch families.

Omega scans are spectrograms produced by the Omicron pipeline from a whitened stream data, offering a visual representation of transient detector noise and gravitational-wave events. A key feature of these representations is their multi-resolution nature: the data — considered in the time-frequency domain — are analyzed simultaneously over multiple quality factors $Q = \frac{f_0}{\Delta f}$, where f_0 is the central frequency and Δf is the bandwidth, enabling sensitivity to transients with different time–frequency morphologies [113–115].

Another important contribution is the iDQ algorithm [116] which autonomously identifies noise artifacts in GW detectors using informations from auxiliary data channels. Interferometers record

thousands of such channels, monitoring potential noise sources including magnetic, environmental or seismic disturbances. The iDQ framework employs wavelet-based features extraction — such as Omicron or SNAX [117] — from these auxiliary channels to generate feature tables and provides a probabilistic assessment of the presence of excess noise.

While noise characterization represents one of the most widespread application area of ML in GW astronomy, interferometric control is indeed another fundamental domain that significantly benefits from such algorithms. In fact, the detector’s behavior is intrinsically non-linear, whereas most control systems are designed under linear assumption. ML algorithms can play a crucial role in approximating these complex dynamics, for instance in the context of lock acquisition and loss-lock prediction — as demonstrated for the LIGO detectors in [118] or [119, 120], which employ Gated Recurrent Unit algorithm (GRU) networks. More recently, [121] introduced a Deep Loop algorithm based on reinforcement learning, which effectively reduces control noise in the low frequency band of LIGO.

In addition, source modeling poses a high-dimensional and computationally demanding challenge. To address this, neural architectures that incorporate physical constraints - such as the Physics-Informed Neural Network [122, 123] and Physics-Informed Neural Operators [124] - have been proposed. Furthermore, ML-based algorithms that accelerate the waveform generation, such as those discussed in [125] and [126] are becoming increasingly relevant.

In GW astronomy, parameter estimation of transient signal plays a central role. Based on Bayesian Inference [127], this task consists in determining the most probable set of source parameters — such as sky location, component masses, and spins — that best describe the observed data. The state-of-the-art parameter estimation framework Bilby [128] remains computationally expensive, requiring substantial resources to perform accurate inference. In this context, ML approaches provide valuable alternatives to accelerate the most computationally intensive phase of Bayesian inference, namely nested sampling. Several promising methods are based on normalizing flows [129], as demonstrated for example in [130]

After this brief introduction in the possible use of ML algorithm in GW astronomy, we focus in enhancing the detection capabilities of a CBC pipeline using machine learning techniques.

4.1 Machine learning for CBC detection

Transient noise, such as glitches, poses a significant challenge in assigning statistical significance to candidate GW triggers. In the ideal Gaussian noise, the SNR of background triggers is expected to remain below $\rho \sim 10$. However, glitches can produce a background distribution of triggers with SNR up to $O(10^2 - 10^3)$, substantially complicating signal discrimination. The ability of CBC pipelines to mitigate the impact of these glitches is therefore crucial in maintaining detection sensibility [131].

A particularly effective method for addressing this issue consists in the time-frequency χ^2 -test described in [78]. In general — for ‘all-sky’ searches, i.e. searches without constraints on event time or sky position — the SNR and χ^2 values are combined algebraically to produce a ranking-statistics capable of more effectively distinguish noise artifacts from the true astrophysical signal.

Building on this idea, one can incorporate multiple pieces of informations, or 'features', to construct a more refined detection statistics, as discussed in Sec. 2.4.1 for the MBTA pipeline. A natural question arises: can a ML algorithm integrate these features more effectively than the current state-of-the-art ranking statistics used in the traditional pipelines?

Several studies have addressed this question such as [132]. In particular, the works [133] for pyCBC and [134] for GstLAL serve as the foundation for our approach. These studies employed multivariate classification techniques, such as Random Forests [135] or Artificial Neural Network [136] to compute novel ranking statistics by incorporating various observables, including SNR, the χ^2 and component masses of the triggers. Furthermore, supervised algorithms have been applied to re-weight the ranking-statistics, and improve the signal-noise discrimination also in burst searches. For example, [137] introduces a decision tree learning algorithm XGBoost and [138] a Gaussian mixture model to classify Coherent WaveBurst triggers, and both methods successfully increase the search sensitivity for short-duration GW in O3 data [139, 140].

While these prior efforts in the CBC context focused on single-detector triggers and explored various input configurations and model architectures, our study extends this approach to coincidence triggers. In fact, we will study the efficiency of the Random Forests algorithm in separating noise from single triggers using the O4a online data acquisition campaign provided by MBTA. In the following sections by considering the O3a and O3b data, we will compare the performance of our machine learning-based ranking statistic against the current MBTA statistic and assess its impact on detection sensitivity taking into account coincidence triggers. Then, we will evaluate the probability of astrophysical origin, namely p_{astro} , associated with the new statistics. We will proceed now by briefly introducing the Random Forests algorithm used in this study.

4.2 Random Forests algorithm

In machine learning literature, we typically consider an input dataset consisting of measured variables — known as features — that determine on one or more target variables, the outputs. The goal is to learn a function that maps inputs to the outputs and then use it to predict outputs for new, unseen input. In supervised learning, each input is associated with a known output label, and the goal is to learn a model that predicts the corresponding outputs for new inputs. Typically an input variable is denoted by the symbol \vec{x} , and single components by x_j . The outputs are denoted as \vec{y} , and they can represent either classes - grouping different inputs - or numerical values. Depending on the nature of the outputs, the problem is categorized as either a classification or a regression task. The difference between the two is that the latter aims to determine the output as a continuous variable, while the former determine whether the output belongs to a class or another. The components x_j of the input are referred to as features. The basic idea behind the learning task is as follows: given an input vector \vec{x} , can we use an algorithm that, by combining the features of \vec{x} , makes an accurate prediction of the output y ? Depending on the task, there are several algorithms that can fulfill this objective. Here, we focus on the Random Forests algorithm, introduced by Breiman [135]. This technique has achieved




Figure 4.1 Left image shows a partition of two-dimensional feature space by recursive binary splitting, as used in CART. Right image shows the corresponding tree. Credits to [9]

great success due to its flexibility and interpretability, particularly in classification problems. This algorithm is derived from the tree-based methods, which is extensively described in [9].

4.2.1 Tree based methods

Tree based methods partition the feature space into a set of rectangles and then fit a simple model — such as a constant — within each region. Well-known algorithms of this type include CART [141] or the C4.5 [142]. Consider for example a simplified problem involving a dataset with two features X_1 and X_2 , taking values in the unit interval, and output y . The feature space can be partitioned using lines parallel to the coordinates axes. In each partition element, we can model y with a different constant. Suppose we initially split the space into two regions and model the response by the value of y within each region. The goal is to choose the best variable and split-point to achieve the best fit. Then, one or both of these subregion may be further split, and this process continues until a stopping rule is applied. In the two-dimensional example we are considering, we first split at $X_1 = t_1$. Then the region $X_1 < t_1$ is split at $X_2 = t_2$ and the region $X_1 > t_1$ is split at $X_1 = t_3$. Finally the region $X_1 > t_3$ is split at $X_2 = t_4$. In this process, we obtain five regions R_1, R_2, \dots, R_5 . In Fig. 4.1 we can see a graphical representation of this procedure.

The classification model then predicts y with a constant value in region R_m according to the input value X_1, X_2 . The tree representation in Fig. 4.1 is particularly instructive. The full dataset sits at the top of the tree. Once the split-point are chosen in each branch of the tree, observations satisfying the condition at each junction are assigned to the left branch, and the others to the right branch. The terminal nodes, or leaves, correspond to the regions R_i . A great advantage of the recursive binary tree is its interpretability, since the feature space partition is fully described by a single tree.

Suppose now we have a p dimensional dataset of N inputs with corresponding response to a set of K classes. So our dataset consists in $\{\vec{x}_i, y_i\}_{i=1}^N$, and $\vec{x}_i = (x_{i1}, x_{i2}, \dots, x_{ip})$ and $y_i \in \{y_1, \dots, y_k, \dots, y_K\}$. Each input belong to a class k and the algorithm is supposed to decide the splitting variables, the

splitting points, and also the structure (topology) the tree should have to predict the correct class. Imagine the features space is partitioned in M regions R_j , we can define the empirical class proportion at node m as

$$\hat{p}_{mk} = \frac{N_{m,k}}{N_m} \quad (4.1)$$

that quantifies the proportion of class k given $N_{m,k}$ as the number of observations in node m belonging to class k and N_m is the total number of observation falling in the node. To decide whether and how to split the node m , we define the node impurity measure Q_m . At each node, the algorithm choses the split that minimizes the impurity — or equivalently maximizes the purity — of the children nodes. There are several expressions for this metrics, the most common used are the Gini index and the Cross-entropy

$$\text{Gini index: } Q_m^{\text{Gini}} = \sum_{k=1}^K \hat{p}_{mk}(1 - \hat{p}_{mk}) \quad (4.2)$$

$$\text{Cross Entropy: } Q_m^{\text{CE}} = - \sum_{k=1}^K \hat{p}_{mk} \log(\hat{p}_{mk}) \quad (4.3)$$

both metrics are 0 for a pure node. Classification at leaf ¹ m to the class $k(m)$ is the majority vote

$$k(m) = \text{argmax}_k \hat{p}_{mk} \quad (4.4)$$

A representation of an example tree can be found in Fig. 4.3. The tree based method shows a significant weak link that is its instability due to the high variance of the trees. A small change in the data can lead to a very different series of splits. This is due to the hierarchical structure of the algorithm, where an error at the top can propagate along the whole structure. A way to reduce this effect is the bagging method, that will be discussed in the following section.

4.2.2 Bootstrap and bagging

Bootstrap is a general technique for assessing statistical accuracy. It is a method for estimating sampling distribution of any quantity obtainable from the original population by re-sampling the data with replacement [143]. The idea behind that is to randomly draw dataset with replacement from the original training data. Each sample has the same number of data as the original one: this procedure is done B times and a Tree is trained on a specific bootstrapped dataset. So if for instance $B = 100$, then we will have 100 datasets build in this way and train 100 trees. Given a training set $Z = (z_1, z_2, \dots, z_N)$, where $z_i = (x_i, y_i)$, the bootstrap dataset is Z^* , the Fig.4.2 represent the bootstrap procedure.

Bootstrap alone is not sufficient to improve the estimate or a prediction by itself. The tool needed in this context is aggregation, or bagging. Consider for instance fitting a model to training dataset

¹A *node* is any internal decision point that may be split; a *leaf* is a terminal node where classification is made.




Figure 4.2 Schematic of the bootstrap process. Datasets Z^* have all the same size N and are drawn from replacement from Z

$Z = (z_1, z_2, \dots, z_N)$ used to fit the predictor $\hat{f}(x)$. Bootstrap aggregation, or bagging, averages this prediction over the models fit to the bootstrapped samples Z^{*b} , $b = 1, 2, \dots, B$, thereby reducing the variance. For each of these Z^{*b} we fit the model obtaining B predictors $\hat{f}_b^*(x)$. The bagging estimate is

$$\hat{f}_{\text{bag}} = \frac{1}{B} \sum_{b=1}^B \hat{f}_b^*(x) \quad (4.5)$$

and it is the average prediction at input x of these B trees. In this way if we consider a regression or classification tree fitted model $\hat{f}(x)$, each bootstrapped tree will typically involve different features from the original dataset and might have different terminal nodes. We can introduce at this point the Random Forests.

4.2.3 Random Forests

Firstly introduced by Breiman in 2001 [135], this algorithm is known for its flexibility, interpretability and solid results. Due to the tree's algorithm structure, applying bagging to these appears quite natural, also, trees can capture complex interaction in the data. Moreover, since each tree output generated in bagging is identically distributed, the expectation of the average of B trees is the same as the expectation of any one of them. Thus, the bias of bagged trees is the same as the one of the individual bootstrapped trees. This means that improvement can be provided only by variance reduction.

The single prediction Tree output $T_b(\vec{x})$ will show a variance $\text{Var}[T_b(\vec{x})] = \sigma^2$ and a pairwise correlation $\rho(\vec{x})$ between distinct trees' prediction. Then, considering an ensemble of such trees, the variance of the average of B such identically distributed random variables will be

$$\text{Var}[T_B(\vec{x})] = \rho \sigma^2 + \frac{1-\rho}{B} \sigma^2 \quad (4.6)$$

It is straightforward that for $B \rightarrow \infty$ the second term disappears, but the first term survives. The idea in Random Forests is to improve the variance reduction of bagging by reducing the correlation between the trees. This is achieved in the tree-growing process through random selection of the input

variables: before each split, we randomly select $m \leq p$ of the input variables as candidates for the splitting. The procedure is explained in the Algorithm 1, with typical value of $m = \sqrt{p}$

Algorithm 1 Random Forests Training Algorithm

- 1: **procedure** RANDOMFORESTTRAIN(D, T, m) \triangleright D : training data, T : number of trees, m : number of features per split
 - 2: **for** $i = 1$ to T **do**
 - 3: $D_i \leftarrow$ BootstrapSample(D)
 - 4: $Tree_i \leftarrow$ TrainDecisionTree(D_i, m)
 - 5: Add $Tree_i$ to forest
 - 6: **end for**
 - 7: **return** forest
 - 8: **end procedure**
 - 9: **procedure** TRAINDECISIONTREE(D, m)
 - 10: **if** stopping criterion is met **then**
 - 11: **return** LeafNode with majority class
 - 12: **end if**
 - 13: $F \leftarrow$ RandomSubsetOfFeatures(m)
 - 14: $f^* \leftarrow$ BestSplitFeature(D, F)
 - 15: $D_{\text{left}}, D_{\text{right}} \leftarrow$ Split(D, f^*)
 - 16: leftSubtree \leftarrow TrainDecisionTree(D_{left}, m)
 - 17: rightSubtree \leftarrow TrainDecisionTree(D_{right}, m)
 - 18: **return** Node(f^* , leftSubtree, rightSubtree)
 - 19: **end procedure**
-

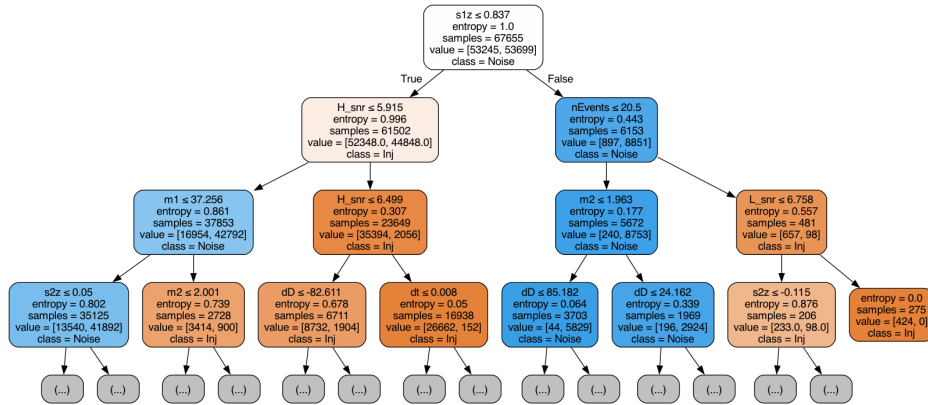


Figure 4.3 Example of tree in a Random Forests used for this study up to three level-depth decision nodes.

Let $T(\vec{x}, \Theta_b)$ be the prediction of the b -th bootstrapped tree at input \vec{x} , and Θ_b is the set of parameters that characterize the tree, such as the split variables, the cut-points and the terminal-node values. A forest made of B such trees will predict for input \vec{x}

$$\hat{f}_{RF}^B = \frac{1}{B} \sum_{b=1}^B T(\vec{x}, \Theta_b) \quad (4.7)$$

In this work we will refer to this quantity as p_s , that is the final score assigned by the forest. For binary classification, $p_s \in [0, 1]$, which represents the estimated probability that the event belongs to a class or the another. This quantity is obtained by averaging the per-tree class probability estimates

$$p_s = \frac{1}{B} \sum_{b=1}^B p_{s,b} \quad (4.8)$$

where $p_{s,b}$ is the leaf probability output by the tree b at input \vec{x} (analogous to $T(\vec{x}, \Theta_b)$). The structure of the Trees is therefore crucial as it determines the terms in Eq. 4.8. The quantities that governs the topology of the tree are known as hyper-parameters, since they are not tuned within the algorithm learning while they are fixed by the user. Changing this hyper-parameters can have a huge impact on the training procedure of the tree and then affecting the classification of the algorithm. In this work we considered the `scikit-learn` package [144], and from now on we will refer to the build-in hyper-parameters provided by it using the `typewriter` font. The set of these hyper-parameters is listed below:

- `n_Estimators`: The number of trees that constitutes the forest. Increasing this number will cause a better performance, since the variance will be reduced. On the other side, a huge amount of trees will cause a consistent computational cost. Also, the performances will not gain precision after a certain values of trees. So choosing the best number of tree is done in order to maximize the precision and minimize the computational time.
- `criterion`: Is the rule used by the tree to split its nodes. the rules usable are `gini` or `entropy`, described above.
- `max_depth`: This hyper-parameter defines the maximum number of levels a single tree is allow to grow in order to obtain the best-purity ending nodes. Each split increase the depth by one increasing the complexity of the three. It has been proved that over-grown trees tend to overfitting, so put a limit in this value is fundamental to get a properly-working classifier.
- `min_samples_leaf`: The minimum number of samples required in each node after a split. A split is at any level of the tree can be considered only if this number of training samples fall in each of the left and right branches.
- `min_samples_split`: The minimum number of samples required at a node to consider splitting it. Nodes with fewer samples than this value becomes leaves.
- `max_features`: The number of features considered to apply the best split. In this work we used only `sqrt`, that means that given a set of m features, each tree will be trained using \sqrt{m} randomly chosen features.

- `ccp_alpha`: Cost-complexity pruning parameter. Higher values prune more of the tree (simpler, lower risk of overfitting); lower values prune less (larger, higher risk of overfitting). The final tree is chosen by balancing training fit with a penalty on tree size controlled by `ccp_alpha`.

These hyper-parameters can span several values, and for each combination of this set, a different tree, and consequently a different classifier can will implemented. The challenge of machine learning then consists in finding the combination of hyper-parameters that provides best algorithm. This procedure is known as 'tuning'. In order to choose this combinations, a metric that compares the different possible mix of parameters is necessary. Several metrics can be used: the most common one is the Accuracy [145], defined as

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (4.9)$$

In a classifier when the output of the algorithm is a label, the TP — True Positive — are the correctly labeled data of the 'positive' class, similarly the TN — True Negative — are the correctly labeled data of 'negative' class. Conversely, the quantities FN — False Negative — and FP — False Positive — represent the wrongly labeled data, respectively of the 'negative' and the 'positive' classes. This metric is not the solely possible choice, another option is given by the F_1 score, defined as the harmonic mean of precision and recall [146], and it is given by

$$F_1 = \frac{2TP}{2TP + FP + FN} \quad (4.10)$$

In general different metrics highlight different aspect. While the Accuracy metric is the 'vanilla' choice, the F_1 metric instead is regarded as the standard metric when both false positives and false negatives are of concern. In this work we explore both of these metrics for tuning the Random Forests algorithm.

4.3 Random Forests for MBTA - single triggers

Applications of multivariate classifier such as Random Forests or Neural Network in the context of CBC search exploiting single detection triggers have been explored by previous works such as [133] and [134] for pyCBC and GstLAL pipelines respectively. In both of these works, the classifiers are trained using single detectors triggers' features — as SNR, χ^2 , masses and so on — and a comparison between the classification score resulting by the classifiers with respect to the state-of-the-art ranking statistic obtained by the pipelines is studied. In particular the work [133] trained the algorithm by considering different combination of features, and noticed that by increasing the number of those, the algorithm where able to outperform the pipeline's performance in noise-signal discrimination.

In this section we study a similar problem, taking into account the single-detection triggers for each interferometers data stream MBTA analyzes. Here, we will explore several combinations of features and, by optimizing the Random Forests architecture maximizing the Accuracy metrics, we

Table 4.1 MBTA triggers - O4a

Triggers	Injections	Noise
H	32396	1499539
L	34703	1747647

will compare the results obtained by the algorithm with respect to the classical MBTA ranking-statistics using the Receiver Operating Characteristic curve [146] for both H and L interferometers using O4a data acquisition campaign.

4.3.1 Dataset, features and hyper-parameters

The LVK collaboration began the first part of the fourth observing run (O4a) on 2023 May 24 at 15:00 UTC and ended it on 2024 January 16 at 16:00 UTC. The dataset considered in this work consists in the single triggers provided by MBTA during O4a data acquisition campaign. For each interferometers, MBTA applies the matched-filtering to the data strain, and every time a trigger shows a $\rho > \rho_{\min}$, a single trigger is recorded. Since we expect few astrophysical signals, while the number of glitch-caused events will be the great majority of the triggers' pipeline, synthetic GW signals are injected² into data. In this way a solid statistical analysis of the detection can be done. In this study, the injected signals are derived from the expected astrophysical population distribution of sources observed in [13, 15]. An Injection is claimed to be 'recovered' if MBTA measures a trigger in the $[-80, 80]$ ms interval centered at the injection time. The injection time corresponds to the peak amplitude of the injected signal, whereas the MBTA trigger time is defined at the end of the matched-filter template. As a result, a systematic offset between injection and trigger times is expected. The chosen coincidence window safely encompasses this offset, accounting for the template duration and timing uncertainties, and therefore provides a conservative criterion for associating triggers with injections while minimizing accidental coincidences. In this way the dataset results in 'injection' and 'noise' labeled signals, where numbers are reported in table 4.1.

The triggers considered here show several features extracted from the MBTA pipeline's analysis. In this study the features considered are:

- The signal-to-noise-ratio SNR ρ for each interferometer, ρ_H for Hanford and ρ_L for Livingston. The threshold on minimum SNR is $\rho_{\min} = 4.4$.
- The autocorrelation-based least-squares statistics computed for both Hanford, $\xi_{PQ,H}^2$, and Livingston $\xi_{PQ,L}^2$.
- The excess trigger rate, used to characterize periods of elevated noise, ER_H, ER_L , as defined in Sec. 2.4.1 (Eq. 2.80).

²injecting a signals means adding the GW signal to the data strain $h(t)$.

- **nEvents:** The number of triggers within the cluster associated to that trigger, as defined in MBTA.
- The component masses (m_1, m_2) and spins (χ_1, χ_2) of the binary system. These parameters are included to evaluate whether specific regions of the parameter space are more susceptible to noise contamination.
- The duration of the waveform template t_{dur} . This choice is motivated by the expectation that many glitches are short-lived compared to astrophysical signals such as BNS. However, high-mass BBH also produce short-duration signals, comparable to glitches.

In Fig. 4.4 we report the distribution for noise and injection population among these features. Since the number of injections and noise events differs, in principle we can possibly work with an imbalanced dataset. The approaches for training and evaluating the algorithm in this context take into account several techniques [147, 148]. Furthermore, since we expect to see few astrophysical events, one can also choose a number of injections that is compatible with the foreseen rate of signals. Even though these ways are feasible, we decided to train and test the algorithm using a balanced dataset: i.e. a dataset where the number of events for each class are equal, or at least are compatible with each others. This choice is not a requirement of the training procedure, but is motivated by the need to control class-dependent biases during training and to ensure that the learned feature representations are equally sensitive to both classes.

Finally, we sampled the events from a randomly shuffled dataset. In this way, we don't risk to select only glitches-affected period of interferometers. Once this sample is done and the balanced dataset is build, this is then split in a training set, that is used to build the decision rules of each tree, according with the hyper-parameters set by the user, and a test dataset, that is needed in order to assess the performance of the trained algorithm. The split is done considering 70% of data for training and 30% for test. In this way we end up with two independent dataset. For the H interferometer we have 22,677 noise labeled triggers, and equal number of 'injection' triggers for training, while for the testing dataset we have 9,719 for 'injections' and 'noise' triggers. For the L training dataset instead we have 24,292 'noise' labeled triggers and the same amount of triggers for 'injections'. For test, 10,411 triggers for the two classes.

At this point we are able to choose the architecture for the Random Forests by choosing the best combination of hyper-parameters that assure a correct prediction of the classes during the training and a robustness in the test phase. To get this combination a grid-search has been done using a k-fold validation: the train dataset is split in k subset and for each hyper-parameters candidate combination and each fold, the model is trained on $k - 1$ folds and evaluated on the remaining fold; validation scores are then averaged across folds. The combination with the highest mean validation accuracy is selected as the final model. The hyper-parameters explored are reported in Table 4.2.

In order to choose the best working architecture, we decided to perform the grid search using all the features described above. In this way we hypothesize that the architecture that classifies best




Figure 4.4 Features pairplot distribution from a subset of 5000 events from dataset. We expect that the Random Forests algorithm is capable in combining these informations of the different population and constructing a statistics that separates the injections from the noise labeled triggers.

the data using all the features, will also performative for less-featured dataset. For the H and L interferometers these values are reported in Table 4.3.

4.3.2 Single triggers results

Once the best hyper-parameters are selected, we trained and test the model using several features combinations. As stated in Eq. 4.8, the output of the Random Forests for a given data point is the ranking p_s , that in this problem we expect it to show values ~ 1 for injection-like triggers, while $p_s \sim 0$ for noise-labelled triggers. At this point we can compare the p_s statistics with the ranking-

Table 4.2 hyper-parameters grid search - O4a

hyper-parameter	values
n_Estimators	15-50-100
criterion	gini-entropy
max-depth	10-25-50
min_samples_leaf	1-5-10
min_samples_split	2-5-10
max_features	sqrt
ccp_alpha	$1 \times 10^{-4} - 5 \times 10^{-4} - 1.5 \times 10^{-3}$

Table 4.3 Best Models - O4a

hyper-parameter	Best Model Values	
	O4a - H	O4a - L
n_Estimators	100	50
criterion	entropy	entropy
max-depth	25	50
min_samples_leaf	1	1
min_samples_split	10	10
ccp_alpha	1×10^{-4}	1×10^{-4}

statistics obtained by MBTA by using the Receiver Operating Characteristic (ROC) curve using the test dataset. This ROC metric consists in defining a false alarm probability α that is the fraction of noise triggers with a classifier score $p_{s,i}$, where i is the index of the trigger, greater than a given threshold \hat{p}_s . Similarly, the Number of Detection N_d is the number of injection triggers with $p_{s,i} > \hat{p}_s$. Defining N_n as the number of noise-labeled triggers and N_s the number of injection-labeled triggers, then these quantities can be defined:

$$\alpha = \frac{1}{N_n} \sum_{i=1}^{N_n} \theta(p_{s,i} - \hat{p}_s) \quad (4.11)$$

$$N_d = \sum_{i=1}^{N_s} \theta(p_{s,i} - \hat{p}_s) \quad (4.12)$$

where θ denotes the Heaviside step function. The ROC curves consist in expressing N_d as function of α . This curve characterizes the trade-off between the detection efficiency and the false alarm probability as the classification threshold \hat{p}_s is varied. By computing Eq. 4.11 and Eq. 4.12 considering the ranking-statistics provided by MBTA — as defined in 2.81 — instead of the ranking p_s , we can perform a comparison between the Random Forests classifier and the traditional pipeline-

Table 4.4 Features sets

Set	Features
F0	ρ, ξ_{PQ}^2
Only Statistical	ρ, ξ_{PQ}^2, ER
F1	$\rho, \xi_{PQ}^2, ER, m_1, m_2, \chi_1^z, \chi_2^z$
F2	$\rho, \xi_{PQ}^2, ER, nEvents, t_{dur}, m_1, m_2, \chi_1^z, \chi_2^z$

based ranking. At this point we trained and test the algorithm considering four different features combination for each H and L model.

The first combination, denoted as *F0*, includes only the signal-to-noise ratio ρ and the ξ_{PQ}^2 test. This minimal feature set allows us to test the ability of the algorithm to assign a ranking to events that reproduces the efficiency of the standard ranking statistics in recovering injections at a given value of α . The next step considers all features related to statistical observables derived from the triggers, adding the *ER* to the set; this configuration is referred to as the *only statistical* dataset. We then extend the feature set to include physical parameters—namely masses, spins, and template duration—together with ρ , *ER*, and ξ_{PQ}^2 , defining the *F1* feature set. Finally, the *F2* set includes all previously listed features. These feature combinations are summarized in Table 4.4.

In Fig. 4.5 are reported the distribution of the ranking-statistics and for the p_s statistics obtained by the algorithm in the *F2* features set.

The ROCs curves relative to the H and L interferometers for each combinations are reported in Fig. 4.6. As we see, just considering the *F0* combination, results in a statistics p_s that is consistent with the ranking-statistics computed by MBTA, although at low α values, that correspond to the most significant triggers, the MBTA statistics is capable in recovering more injection with respect to p_s . Adding the *ER* in the 'Only Statistical' configuration, we obtain curves for p_s that are slightly better than the ranking-statistics one for all α values, suggesting that the relations found by MBTA for these features are sub-optimal. If we consider the physical features such as the masses and the spins in *F1*, or even the template duration t_{dur} and the *nEvents* in *F2* results in statistics p_s that considerably outperform the ranking-statistics, suggesting that including these features can highly benefit the injections recovery and, consequently, the real events triggering.

Since the features set *F2* shows a ROC curve from a p_s that significantly surpasses the MBTA ranking-statistics, it is interesting to see how the Random Forests statistics is capable in separating noise from signals at a given α value in the physical parameters space. In order to do so, we divide the total mass M_{tot} and the effective spin $\chi_{eff} = \frac{m_1\chi_1^z + m_2\chi_2^z}{m_1 + m_2}$ space in 10 intervals and compute the ratio between the recovered injections for both p_s and the ranking statistics. More formally, we define the ratio *R* as

$$R = \frac{N_d(\bar{p}_s)}{N_d(\bar{\rho}_{rw})} \Big|_{\bar{\alpha}} \quad (4.13)$$

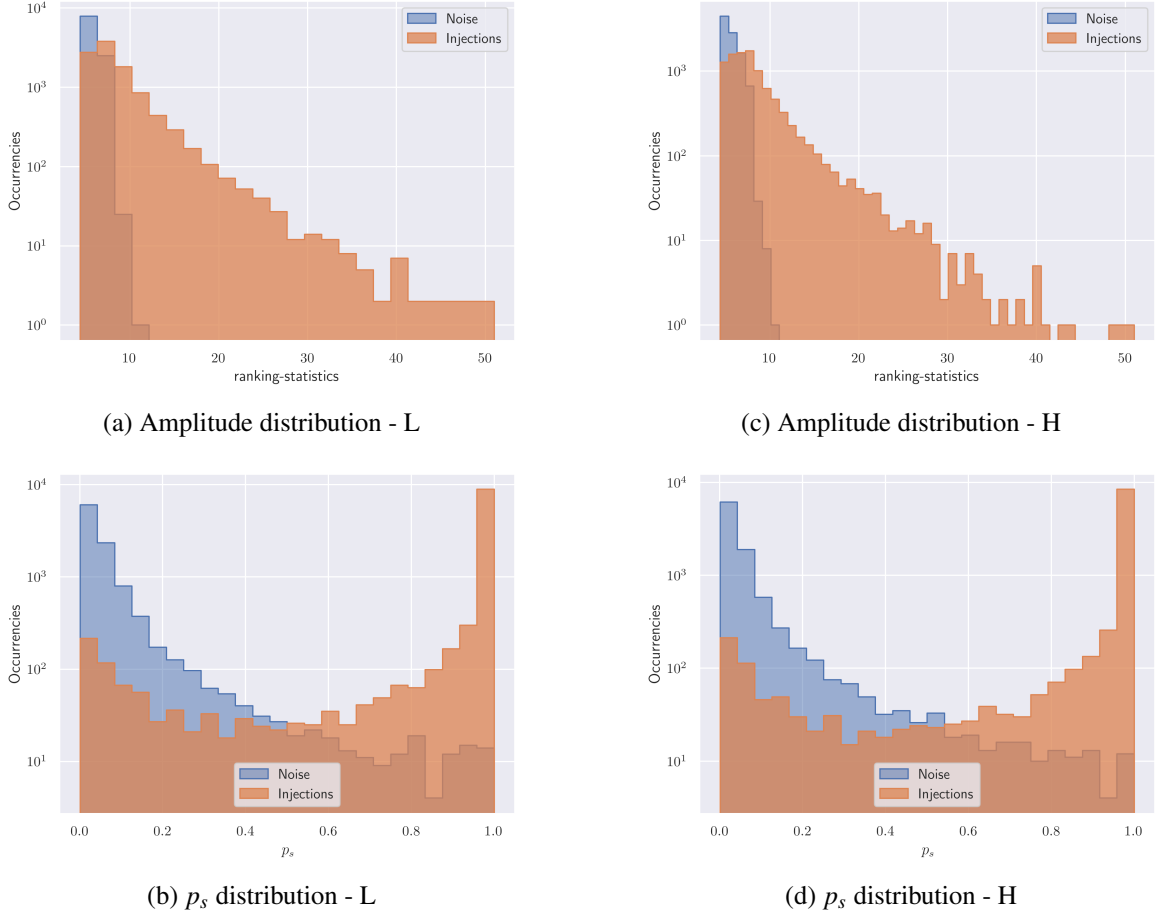



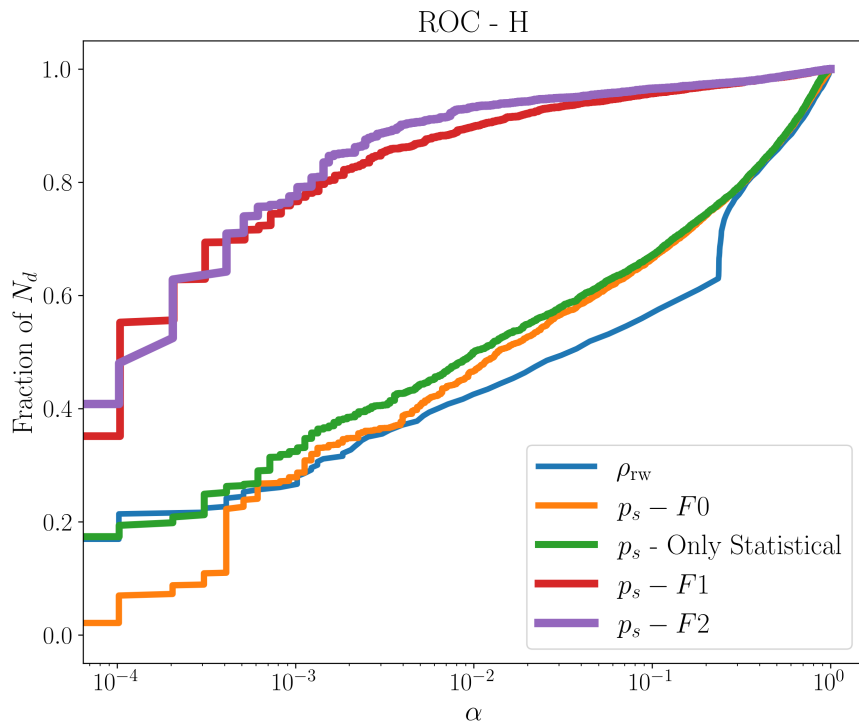
Figure 4.5 Distributions of ranking-statistics (amplitude) and p_s for L and H detectors. The top plots show ranking-statistics distributions; the bottom plots show p_s distributions for the same test dataset.

where $\bar{\alpha}$ is a fixed value of α that corresponds to two specific values of $\bar{\rho}_{rw}$ and \bar{p}_s . N_d is the number of recovered injections, at those fixed values for the corresponding considered statistics. The ratio R quantifies then how well the p_s statistics recovers injections in that physical interval. The $\bar{\alpha}$ values chosen as threshold is $\bar{\alpha} = 10^{-2}$ that correspond to $\bar{\rho}_{rw} = 8$ for both interferometers. The corresponding \bar{p}_s values are $\bar{p}_s = 0.7$ for H, while $\bar{p}_s = 0.65$ for L. The results are reported using logscale in Fig. 4.7.

From these figure we see that in all physical space the injections are recovered with greater efficiency using p_s with respect to the ranking-statistics results. The solely areas where the R is in favor of MBTA statistics are the lower χ_{eff} - low M_{tot} areas and high χ_{eff} - low M_{tot} area for H interferometer. One reason for this behavior may be the relative lack of injections in such parameters' area (see Fig. 4.8), yet this explanation is incomplete, since the L case, which is analogous to the H in terms of injections, does not present this problem. Further investigations are required.



(a) ROCs curve for O4a - L



(b) ROCs curve for O4a - H

Figure 4.6 Receiver Operating Characteristic (ROC) curves for the best models of O4a.



(a) $M_{\text{tot}} - \chi_{\text{eff}}$ R distribution among intervals - H(b) $M_{\text{tot}} - \chi_{\text{eff}}$ R distribution among intervals - L

Figure 4.7 Efficiency in injection recovery ratio between p_s and ρ_{rw} for F_2 features set. Events binned in M_{tot} and χ_{eff} . Each box shows $\log_{10}(R)$, where $R = \frac{N_d(\bar{p}_s)}{N_d(\bar{\rho}_{\text{rw}})} \Big|_{\bar{\alpha}}$. Green gradient indicates $R > 0$ (higher recovery with p_s), conversely red gradient is for $R < 0$ (higher recovery with ρ_{rw}) and white bins are where no injections at all were performed.



(a) $M_{\text{tot}} - \chi_{\text{eff}}$ injections - noise distribution among intervals - H(b) $M_{\text{tot}} - \chi_{\text{eff}}$ injections - noise distribution among intervals - L

Figure 4.8 Injection and noise population for H and L. The net density of injection of the upper left corner of the plot with respect to the bottom may give some hints about the lack efficiency in detection for H, yet we do not appreciate a significant difference between the IFOs regarding the population distribution, this may suggest that other causes affects this H mis-recovery.

4.4 Random Forests for MBTA - coincidence triggers

In the previous section, we studied a possible application of a Random Forests algorithm to MBTA single-detector triggers considering the O4a data acquisition campaign. Here we push forward the study and try to assess the astrophysical origin of a trigger, named p_{astro} , considering the double-coincidence triggers obtained during the O3 data acquisition campaign via the p_s statistics, testing the consistency of this method by comparing the results with the catalogues' events p_{astro} .

The dataset consists of the triggers produced by MBTA during the O3a sub-run (1 April 2019 15:00 UTC - 1 October 2019 15:00 UTC) and O3b (1 November 2019 15:00 UTC - 27 March 2020 17:00 UTC). We consider the HL-Von coincidence triggers, defined as those produced by MBTA when all three interferometers — Hanford (H), Livingston (L) and Virgo (V) — were simultaneously operational (i.e. locked and acquiring data). A coincidence trigger is claimed when both H and L triggers satisfy the condition $\rho_H, \rho_L > \rho_{\text{min}} = 4.5$, and the time difference between the two falls within the predefined coincidence window. Also ρ_V is recorded for these events. The injections are claimed to be recovered if a pipeline trigger occurs within the window $[-80, +40]$ ms of the injection time, similarly as done previously and in conformity with the O3 injection recovery policies. The asymmetric window reflects the fact that the injection time corresponds to the peak amplitude of the signal, whereas the pipeline trigger time is defined at the end of the matched-filter template. In addition, since the injection time is the time of the maximum amplitude, while the recovered time is the end of the template, i.e. with the ringdown for BBH, a correction of $100(\frac{m_1}{180})$ ms, where m_1 is the primary mass, is applied when making the time difference between the injected and recovered time. Among all possible triggers that fall in such windows, the loudest one is picked. The injections used are those performed during the O3 observing campaign, and are obtained from a population model designed to approximate the expected distribution of true astrophysical sources, as described in [13] and [15]. The datasets' numbers are reported in Table 4.5.

To build the training and test dataset, we followed the same philosophy as the single-detectors case: we used a balanced dataset containing the same number of noise and injections for each sub-set, that in this case correspond in O3a to 58,281 noise triggers and an equal number for injections in the training dataset and 24,977 triggers for noise and 24,977 for injections in the test dataset. In O3b we got 53,472 noise triggers, with same number of injections for the training set, and 22,917 triggers for each class in the test.

The sampling procedure is the same as the one of the previous section, ending up in this case too with two independent datasets, each with balanced representation of population. The features

Table 4.5 O3 datasets

sub-run	Noise	Injections	Astrophysical Events
O3a	173130	83258	23
O3b	129133	76389	16

Table 4.6 hyper-parameters grid search- O3

hyperparameter	values
n_Estimators	15-50-100
criterion	gini-entropy
max-depth	10-12-15
min_samples_leaf	1-5-10
min_samples_split	2-5-10
ccp_alpha	None, 1×10^{-4} – 5×10^{-4} – 1.5×10^{-3}

used are the SNR for HL-Von coincidence triggers ρ_L, ρ_H , each requiring to satisfy $\rho > \rho_{\min}$. The autocorrelation-based least-squares statistics for both Livingston and Hanford $\xi_{\text{PQ,L}}^2$ and $\xi_{\text{PQ,H}}^2$, along with the Excess Rates ER_L and ER_H . As done previously, we consider also the masses (m_1, m_2) and the spins (χ_1^2, χ_2^2) of the template that triggered. Analogously, t_{dur} and nEvents are included in the training. Finally, we used the difference in time Δt , in phase $\Delta\phi$ and in distance ΔD evaluated between H and L triggers. These features, incorporating the informations of the double coincidence, capture inter-detector consistency and are included to allow a classifier to learn pattern associated with true coincidence events versus noise. In fact, triggers coming from astrophysical events are expected to show correlations across detectors, not just between their arrival times but also regarding the phase and amplitude of the signal [149].

The best hyper-parameters set have been chosen by applying a grid-search among several possible values, reported in Table 4.6. We highlight that in this case we imposed a greater constriction to the depth of the trees, forcing this value to a maximum of 15. This has been done since we found out empirically that in this way the model was much less prone to overfitting. The metrics used in this case is the $F1$ metric. The final hyper-parameters combinations are reported in Table 4.7. The architecture is similar to the single-detector design, but few hyper-parameters indicates that the algorithm must adapt to a different type of dataset. In fact, the `ccp_alpha` parameter is taken to zero, while less limitations are imposed to the `min_samples_split`, allowing a more free growth of the trees.

Table 4.7 Best Models - O3

hyperparameter	Best Model Values	
	O3a	O3b
n_Estimators	100	100
criterion	entropy	entropy
max-depth	12	15
min_samples_leaf	1	1
min_samples_split	2	5
ccp_alpha	None	None

The Random Forests using these configurations for the corresponding dataset is then trained and tested. From Fig. 4.9, in particular the top panels showing the ranking-statistics distribution for the two population, we see that imposing the coincidence is an efficient way to significantly reduce the noise impact. The tails of the distribution for noise-labeled triggers is shorter than the one experienced for single detector trigger. As previously, the output of the algorithm, the rank p_s , is compared against the ranking-statistics obtained by MBTA via ROC curves. We expect then that the p_s statistics will not improve significantly the detection efficiency in this dataset with respect to the ranking-statistics.

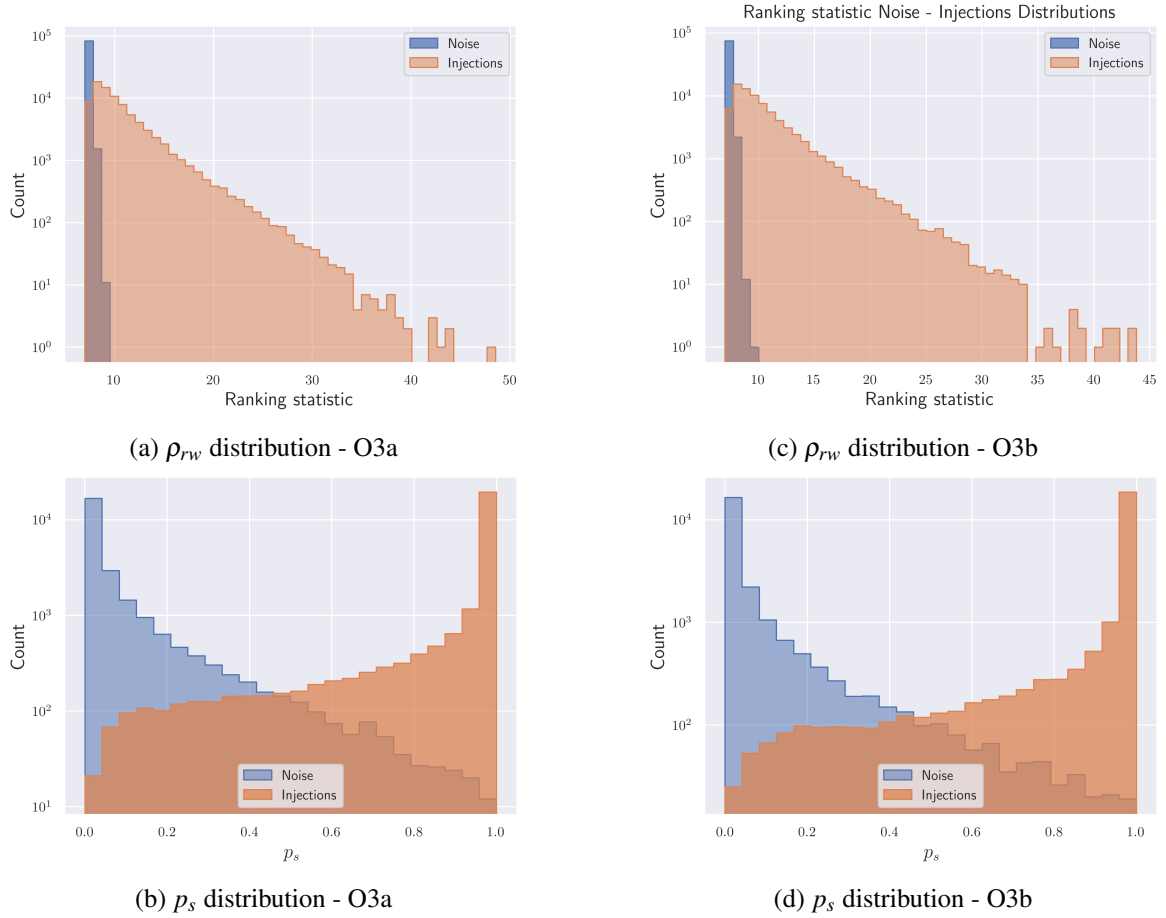




Figure 4.9 Distributions of ranking-statistics (p_{rw}) and p_s for O3a and O3b HL-Von coincidence MBTA triggers. The top plots show ranking-statistics distributions; the bottom plots show p_s distributions for the same test dataset.

The ROC curves, computed with the same procedure as done for the single detector case, are reported in Fig. 4.10. We see that the efficiency in injections recovery is really similar to the ranking-statistics obtained by MBTA. In fact, even if the p_s distribution is better than Fig. 4.5, yet the p_{rw} for double coincidence is extremely good in separating signal from noise.

The Random Forests rank p_s in this case is not outperforming the MBTA ranking-statistics, yet it is understandable since by imposing the coincidence we are able to remove a consistent part of the



(a) ROC curve for O3a



(b) ROC curve for O3b

Figure 4.10 Receiver Operating Characteristic (ROC) curves for HL-Von coincidence triggers (O3). Even though the red curve obtained from the p_s statistics is almost everywhere above the ranking-statistics one, the difference is not substantial as we saw in the single detector case.

background. So, the results are compatible with the state-of-the-art statistics. Let's now compute the p_{astro} for the test dataset given this new rank.

4.5 Random Forests for MBTA - p_{astro} computation

As introduced in Chapter 2 (Eq. 2.85), p_{astro} is the probability that a given trigger is of astrophysical origin. We can use the statistics p_s to build the distributions $R(p_s)_S$ and $R(p_s)_N$. These rate densities can in fact be factorized as

$$R(p_s)_S = p(p_s|s)\Lambda_1 \quad (4.14)$$

$$R(p_s)_N = p(p_s|n)\Lambda_0 \quad (4.15)$$

where we considered the effective rate, that is the expected number of astrophysical events during the data acquisition campaign's time, Λ_1 , or the effective number of noise triggers Λ_0 . Each of these rates are multiplied for the probability density functions $p(p_s|s)$ for signals and $p(p_s|n)$ for noise. The p_{astro} results then in

$$p_{\text{astro}} = \frac{p(p_s|s)\Lambda_1}{p(p_s|s)\Lambda_1 + p(p_s|n)\Lambda_0} \quad (4.16)$$

The astrophysical rates Λ_1 considered in this study are the number of astrophysical events reported in the catalogues GWTC-2 (O3a) and GWTC-3 (O3b). These are the events with an inverse false alarm rate $\text{IFAR} > 0.5\text{y}$ ³, as reported from the LIGO-Virgo-Kagra collaboration during the O3 campaign [13, 15]. Then the values are $\Lambda_{1,\text{O3a}} = 39$ for O3a and $\Lambda_{1,\text{O3b}} = 35$ for O3b. The choice of these values is conservative since it takes into account only the significant events reported during the acquisition campaign, and ignore all the possible sub-thresholds signals — i. e. real but not significant events that do not survive the catalogue threshold cut. Other possible values of the astrophysical rates, as suggested in [150] and [151] is to estimate the signal rate prior with self-consistent interface, using a likelihood from the observed distribution.

Once the Λ_1 values are set, we need to estimate the probability density functions $p(p_s|s)$ and $p(p_s|n)$. Since no theoretical models are known for modeling the output of a classifier, we relied on Kernel Density Estimation. This technique introduced by Rosenblatt [152] consists in extrapolating a probability density function $\hat{f}(x)$ given a set of identically distributed random variables $\{x_i\}_{i=1}^n$ on n points that follow the unknown distribution $f(x)$. This technique relies on assigning to each point x_i a function $K(x_i, t)$ that gives to the corresponding position of the input an interval. This function is known as kernel function. Here t represents the point at which the estimated density $\hat{f}(t)$ is evaluated, while x_i are the observed data values. The approximated $\hat{f}(t)$ function is obtained then

$$\hat{f}(t) = \frac{1}{n} \sum_{i=1}^n K(x_i, t) \quad (4.17)$$

³The IFAR quantifies how often a noise trigger would produce a ranking statistic equal to or higher than that of the candidate event. An $\text{IFAR} > 0.5, \text{yr}$ indicates that a similar noise event would occur, on average, less than once every six months.

The kernel function satisfies the properties of being a positive-defined function ($0 \leq K(x, t) \leq \infty$) and the normalization $\forall x, \int_{-\infty}^{\infty} K(x, t) dt = 1$. This last condition ensures the normalization of the approximated $\hat{f}(t)$. Mostly, symmetric functions are used for this task, with expression

$$K_{\text{sym}} = K(x, t) = \frac{1}{h} K\left(\frac{x-t}{h}\right) \quad (4.18)$$

The quantity h is the bandwidth of the kernel function, and it is a parameter that must be wisely used since it defines the smoothness of the resulting pdf. A too big h value will end into a too smooth function, leading into a loss in informations about locally-relevant behavior, on the other side, a too small bandwidth will cause a meaningless distribution, since it will follows all the statistical oscillation of the data. In fact, as n increases and h decreases appropriately, $\hat{f}(x)$ is expected to converge toward the true underlying distribution $f(x)$. In this work, we used a Gaussian Kernel function with empirically defined bandwidths for noise and injections pdfs, and are $h_s = 0.8$ for signal distribution and $h_n = 0.6$ for noise. The $p(p_s|s)$ and $p(p_s|n)$ are then approximated using this method applied to the p_s distributions obtained by the Random Forests for both the data acquisition campaigns, albeit a transformation to the statistics p_s was required since it takes values only in $[0, 1]$, and border effect can invalidate the KDE. So, a logit transformation is applied to the data as described in Eq. 4.19.

$$\tilde{p}_s = \ln\left(\frac{p_s}{1-p_s}\right) \quad (4.19)$$

The result of this transformation \tilde{p}_s is then the input of the KDE that gives us the pdfs required to compute the p_{astro} . The distribution of \tilde{p}_s along with the KDE result for O3a is reported in Fig. 4.11.

With these distributions, along with the Λ values defined above, the computation of p_{astro} for noise-and injection labelled populations is possible. The results for the test dataset are given in Fig. 4.12 where we plot the p_{astro} obtained for the two population with respect to the ranking-statistics. This has been done in order to check the consistency between the p_s and ρ_{rw} in computing the astrophysical probability.

At this point, we can test the p_{astro} values obtained via p_s statistics for the events reported by MBTA in GWTC-2 and GWTC-3, that are, respectively, O3a and O3b data acquisition campaign. For this study, the events reported in the catalogues have been previously removed from the dataset, then, once the model is trained and the PDFs are build using the test dataset, the $p_{\text{astro}}^{(p_s)}$ are computed for these events — the p_s superscript indicates the statistics which have been used to compute the astrophysical-source probability. The MBTA events consist in 23 coincidence events for O3a and 16 for O3b. Among the O3b events, 5 show a $p_{\text{astro}} < 0.5$ and have been maintained to check if the Random Forests assign those a p_s such that $p_{\text{astro}}^{(p_s)}$ may improve. The results are depicted in Fig. 4.13.

We see that compatible results are obtained for all the events with ranking statistics above 10 in both datasets, except for GW190924_021846 event in O3a (Fig. 4.13a). In O3a, for ranking statistics below 10, there are three events with $p_{\text{astro}}^{(p_s)} < 0.5$, and only one of them has a ranking statistics greater than 10—namely, the GW190924_021846 event. This case will be discussed in detail. In O3b (Fig.




Figure 4.11 Probability density function obtained for O3a \tilde{p}_s distributions for signal (red) and (blue) labeled triggers.

4.13b), for ranking statistics below 10, there are five events with $p_{\text{astro}}^{(p_s)} < 0.5$ with no anomalous behavior for events with ranking statistics above 10.

The event GW190924_021846 is reported as a vanilla binary black hole (BBH) merger, with source-frame component masses reported by GWOSC [153] as $m_1 = 8.9_{-2.0}^{+7.0} M_\odot$ and $m_2 = 5.0_{-1.9}^{+1.4} M_\odot$. The corresponding MBTA trigger does not exhibit any apparent anomalies, and the online parameters provided by the pipeline are $m_1 = 41.3 M_\odot$, $m_2 = 1.97 M_\odot$, with a ranking statistic value of 10.9. This discrepancy is expected because the masses reported by the MBTA pipeline correspond to the parameters of the template that maximizes the matched-filter signal-to-noise ratio, rather than to astrophysical source parameters inferred from Bayesian parameter estimation. These template masses are detector-frame quantities drawn from a discrete template bank and are optimized for detection efficiency rather than physical accuracy. As a result, significant differences between online trigger parameters and the source-frame masses reported by GWOSC are not uncommon. However, the classifier unexpectedly fails to assign a high score to this event. Specifically, the value of $p_{\text{astro}}^{(p_s)}$ is only 0.04, which is inconsistent with the event’s amplitude and significance. Upon further inspection, it was empirically observed that removing the ER_H and ER_L features from the input feature set led to a substantial improvement in the classification of this event. As shown in Fig. 4.14, the $p_{\text{astro}}^{(p_s)}$ score shifts from 0.04 to $p_{\text{astro}}^{(p_s, \text{noER})} = 0.98$. This dramatic change suggests that the excess rate features (ER_H , ER_L) heavily bias the classifier against this particular event. The reasons for this behavior are not yet understood and require further investigations, especially given that the event GW190924_021846 otherwise exhibits high significance and physical plausibility.





(a) p_{astro} - O3a - injections (orange), noise (red)(b) p_{astro} - O3b - injections (orange), noise (red)

Figure 4.12 p_{astro} values obtained via Kerner Density Estimation over test dataset considering the \tilde{p}_s statistics for O3a (left) and O3b (right) triggers. We see that all the noise-labelled signal lay in the bottom part of the plot, with no triggers above 0.15.

For completeness, in Tables 4.8 and 4.9 all the, p_{astro} and $p_{\text{astro}}^{(p_s, \text{noER})}$, that are the $p_{\text{astro}}^{(p_s)}$ values calculated without the *ER* features, are reported for all the GWTC-2 and GWTC-3 HL-Von MBTA events. We highlight that no events with ranking statistics above 10 show a worse behavior in $p_{\text{astro}}^{p_s}$



(a) $p_{\text{astro}}^{(p_s)}$ and p_{astro} - O3a MBTA catalogue's events



(b) $p_{\text{astro}}^{(p_s)}$ and p_{astro} - O3b MBTA catalogue's events

Figure 4.13 Comparison between $p_{\text{astro}}^{(p_s)}$ values obtained by p_s statistics and p_{astro} for O3 MBTA catalogues' events.

— whether with or without the ER — with respect to the catalogue's results, thus we decided not to discuss them, since their significance is relative.

This method shows consistency with the results obtained by the standard MBTA pipeline for triggers that show a ranking statistics above 10. So, we decided to perform an agnostic search for new

Table 4.8 O3a Events p_{astro} comparison

Event name	ranking-statistics	p_{astro}	$p_{\text{astro}}^{(p_s, \text{ER})}$	$p_{\text{astro}}^{(p_s, \text{noER})}$
GW190408_181802	14.05	1.00	1.000	0.999
GW190412	17.62	1.00	1.000	0.999
GW190413_134308	9.59	0.99	1.000	0.999
GW190503_185404	11.36	1.00	1.000	0.999
GW190512_180714	11.28	0.99	1.000	0.999
GW190513_205428	10.38	0.99	1.000	0.999
GW190517_055101	10.37	1.00	1.000	0.999
GW190519_153544	13.26	1.00	1.000	0.999
GW190521	10.78	0.96	1.000	0.999
GW190602_175927	11.91	1.00	1.000	0.999
GW190701_203306	7.74	0.87	0.036	0.055
GW190706_222641	11.31	1.00	1.000	0.999
GW190720_000836	10.56	1.00	1.000	0.999
GW190725_174728	9.03	0.59	0.775	0.090
GW190728_064510	12.86	1.00	1.000	0.999
GW190803_022701	8.57	0.96	1.000	0.999
GW190828_063405	14.85	1.00	1.000	0.999
GW190828_065509	10.00	0.96	1.000	0.999
GW190915_235702	12.35	1.00	1.000	0.999
GW190916_200658	7.94	0.66	0.076	0.070
GW190929_012149	9.21	0.64	1.000	0.999
GW190727_060333	8.72	1.00	0.090	0.042
GW190924_021846	10.89	0.99	0.037	0.978

Table 4.9 O3b Events p_{astro} comparison

Event name	ranking-statistics	p_{astro}	$p_{\text{astro}}^{(p_s, \text{ER})}$	$p_{\text{astro}}^{(p_s, \text{noER})}$
GW191105_143521	9.89	1.00	1.000	1.000
GW191113_071753	8.73	0.68	0.063	1.000
GW191127_050227	9.18	0.73	1.000	1.000
GW191215_223052	10.12	0.99	1.000	1.000
GW191230_180458	8.85	0.40	1.000	1.000
GW200115_042309	10.56	1.00	1.000	1.000
GW200208_130117	9.43	1.00	1.000	1.000
GW200208_222617	8.25	0.02	0.011	0.007
GW200209_085452	8.75	0.97	1.000	1.000
GW200216_220804	8.01	0.02	0.031	0.060
GW200219_094415	9.88	1.00	1.000	1.000
GW200224_222234	18.07	1.00	1.000	1.000
GW200308_173609	7.47	0.24	0.000	0.000
GW200311_115853	14.55	1.00	1.000	1.000
GW200316_215756	8.77	0.03	1.000	1.000
GW200322_091133	8.26	0.62	0.001	0.001

Chapter 5


Convolutional Neural Network for gravitational-wave triggering

In the previous chapter we studied a possible application of Random Forests algorithm for post-processing triggers provided by the MBTA pipeline. Instead here we focus on developing a machine learning algorithm for triggering GW candidates from interferometers' data strain. Detecting GW signals using machine learning driven approach is a possibility explored by several studies. In this field a number of pipelines have been developed and deployed for triggering possible GW candidates. For CBC searches algorithms such as Aframe [155] or AresGW [156] shows application of Residual Network architecture that are capable in detecting signals at sensitivities comparable to the classical matched-filtering pipelines — at least in specific area of the parameter space — at a fraction of cost. Along with these, also a Deep Learning pipeline for bursts signals, the Mly pipeline [157] has been deployed for detecting un-modeled signals. These pipelines put a great effort in studying the best architecture achievable for this kind of problem, along with exploring possible training procedure to maximize the algorithms' performances.

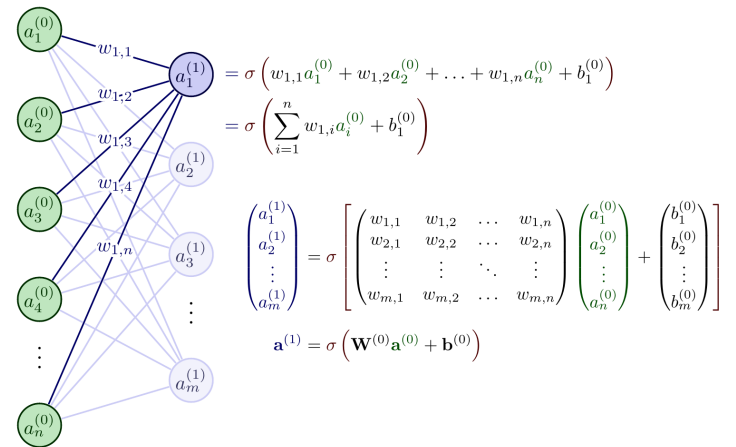
Such works — apart of being effective scientific-deployed pipelines — show how promising is this research field, and this study follows the path of these previous analysis, exploring a possible application of Deep Learning applied to GW candidates triggers. In this section we investigate a proof-of-concept Convolutional-Neural-Network (CNN) powered detection algorithm for GW data analysis. Following an idea developed in [158], we construct features-rich data from time-series SNR given by a matched-filtering analysis. This procedure results then in images that can be used to train and test the CNN, exploiting in this way the advantages this algorithm developed in the years for this kind of problems. Considering then several simulation campaigns varying noise assumption and astrophysical sources properties and parameters space, we finally show the results comparing the performance obtained by such algorithm with respect to a classical matched-filtering search.

5.1 Convolutional Neural Network

We start by briefly reviewing the Convolutional Neural Network [159] architecture, from the basic neuron up to the modern Residual Network block. In its seminal article 'The perceptron: a probabilistic model for information storage and organization in the brain' [160], F. Rosenblatt described how mathematical function can resemble functionalities similar to human neuron by relating inputs to an activation function that 'fires' when specific value of the inputs are passed. A scheme of such structure is shown in Fig. 5.1a.



(a) Scheme of a perceptron. The inputs x_i are passed to the activation function after being multiplied by weights w_j and summed. If the result is above a specific activation threshold, the perceptron activates and gives an output.



(b) Pictorial representation of a neural network layer. The inputs are provided from the previous layer — the light green dots — and are connected to the following neuron layer — the light violet dots — with a specific weight w_{ji} , where i and j are, respectively, the indices of the input and output neurons. In each output neuron an activation function σ takes as input the weights and outputs of the previous layer neurons and produces an output.

Figure 5.1 Illustrations of (a) the perceptron structure and (b) the mathematical formalism behind neural network layers.

The single perceptron, or neuron, can be activated given a set of inputs (x_i) and weights (w_j) passed to the activation function $\sigma(x_i, w_j)$. Collecting several neurons and stacking them in a layer is the approach developed to collect informations from different neurons, each with its own activation function and set of weights associated to the input. A given input then correspond for each neuron to a weighted sum and a corresponding activation function. By inserting one layer after the other, the output of the previous neuron layers becomes the input of the following, with an activation procedure similar to the one just mentioned. This structure is depicted in Fig. 5.1b.

By inserting recursively layers, possibly varying the number of neurons, one obtain what is called a Neural Network (NN) [161], see Fig 5.2. This architecture has been the first Deep Learning algorithm developed for tasks such as data labelling or image recognition. The whole point about




Figure 5.2 Example of a Neural Network with 5 layers. The green light dots represent the input, the light violet are the hidden layers and the light red one is the output layer.

this algorithm is that the NN can approximate non-linear behavior of data in their parameter space, something that the Random Forest or other algorithms such as the K-cluster cannot do. This advantage came with the price of data-hunger of such algorithms, that requires impressive amount of data to properly work. Now we can ask ourself, how do such algorithms learn? What is the procedure to assign the correct weight and activation threshold to each neuron? In addition, the learning-algorithm must be fast, since usually we deal with millions or even billions of parameters.

The technique developed for this kind of task is the back-propagation algorithm along with the stochastic gradient descending. We aim to find a configuration of weights that ensure that for each input vector, the output produced by the network is sufficiently close to the desired result. Following [162] and referring to Fig. 5.1b, each neuron in the $(l + 1)$ -th layer receives inputs from all neurons in the previous layer l . Denoting the activations of layer l as $a_i^{(l)} = y_i$ and the weights connecting neuron i in layer l to neuron j in layer $l + 1$ as $w_{ji}^{(l)}$, the total input to neuron j in layer $l + 1$ is

$$x_j^{(l+1)} = \sum_i w_{ji}^{(l)} y_i^{(l)}, \quad (5.1)$$

The activation of neuron j is then obtained by applying a nonlinear activation function σ :

$$y_j^{(l+1)} = \sigma(x_j^{(l+1)}), \quad (5.2)$$

where, in the following, the logistic (sigmoid) function is used:

$$y_j = \sigma(x_j) = \frac{1}{1 + e^{-x_j}}. \quad (5.3)$$

but any activation function which has a bounded derivative will work. At this point we can create an error function E , i.e. a function that quantifies how distant are the results obtained by the network with respect to the attempted one:

$$E = \frac{1}{2} \sum_c \sum_j (y_{j,c} - d_{j,c})^2 \quad (5.4)$$

where c is the index that identifies the cases of the problem, for instance is the index class corresponding to a picture in a image-recognition problem and j indexes over the final output units ($a_j^{(L)}$) of layer L related to the $L - 1$ layer. From now on we will suppress the layer-indices to slim the notation. The quantities $d_{j,c}$ is the desired output, while $y_{j,c}$ is the result obtained by the network. The variation of the error function, considering only a particular case of c and suppressing the index, for the specific output y_j , is

$$\frac{\partial E}{\partial y_j} = y_j - d_j \quad (5.5)$$

We can then compute the change in input x_j that affects E as

$$\frac{\partial E}{\partial x_j} = \frac{\partial E}{\partial y_j} \frac{\partial y_j}{\partial x_j} = \frac{\partial E}{\partial y_j} y_j (1 - y_j) \quad (5.6)$$

where we have used the derivative of the sigmoid activation function $\sigma'(x_j) = y_j(1 - y_j)$. This means that we know how a change in the total input will affect the error. Since the total input x_j is a linear combination of the weights w_{ji} and the activation y_i of the lower level units, we can express how the error will be affected by changing their states and weights. The derivative of the error E with respect to to the weight w_{ji} is

$$\frac{\partial E}{\partial w_{ji}} = \frac{\partial E}{\partial x_j} \frac{\partial x_j}{\partial w_{ji}} = \frac{\partial E}{\partial x_j} \cdot y_i \quad (5.7)$$

and regarding the output of the i -th unit the contribution to $\partial E / \partial y_i$ due to the effect of i on j is

$$\frac{\partial E}{\partial x_j} \frac{\partial x_j}{\partial y_i} = \frac{\partial E}{\partial x_j} \cdot w_{ji} \quad (5.8)$$

by summing on the j index we take into account all the connection emanating from unit i as

$$\frac{\partial E}{\partial y_i} = \sum_j \frac{\partial E}{\partial x_j} \cdot w_{ji} \quad (5.9)$$

so we can compute the $\partial E / \partial y$ for any unit in the penultimate layer when given $\partial E / \partial y$ for all units in the last layer. This procedure can be repeated for the successively early layers, computing $\partial E / \partial w$ and changing the weight accordingly after every input-output case by minimizing the error

function. The simplest version is to implement the gradient descent by changing each weight by an amount proportional, let's say by a constant ε , to $\partial E / \partial w$:

$$\Delta w = -\varepsilon \partial E / \partial w \quad (5.10)$$

This method has the advantage of simple implementation and parallel computation and it can be further fasten-up by considering the stochastic-gradient-descent approach [163]. This approach is the corner-stone method for Deep Learning technique. More advanced architectures such as Autoencoder or Convolutional Neural Network (CNN) are based on this neuron connections and backpropagation optimization.

The CNN in particular are models specifically developed for images recognition. The main idea behind this algorithm is to extract from images shapes and features that characterize the objects inside the picture and at the same time not to be sensitive to the position inside the image. The operation that is suitable for these goals, that are features extraction and space invariance, is the convolution, that is implemented in the convolutional layers. Consider a kernel K , known as filter, consisting in a matrix ($M \times N$) — clearly the numbers M and N defines the size of the filter — and an input I of ($A \times B$) pixels (or features). The convolutional operation consist in multiplying the values of the filter for the corresponding pixels of the image, and then summing the result of these products. More formally, the output of of the convolution $C(i, j)$ is


$$C(i, j) = (I \star K)(i, j) = \sum_m^M \sum_n^N I(i+m, j+n) \cdot K(m, n) \quad (5.11)$$

where i, j and m, n are, respectively, the features and kernel coordinates. We see that convolution consists in the sum of the element-wise product of the filter K and a portion of the input. The number of pixel we move the filter across the image is known as stride; in this way, with strides bigger than 1, we can scan the image by jumping an amount of pixels equal to the chosen stride value and obtaining in this way a sparser coverage of the input feature map. Large strides are useful since they increase the field of view of the filter, capturing global features of the input. Also, the stride operation reduce the dimension of the input at the following step resulting in a reduction of the computational load and memory usage. In fact, the dimensions W (width) and H (height) of the output features after a convolutional operation with filter size F — that is a square kernel K of size ($F \times F$) — with stride S is

$$W_{\text{out}} = \frac{W_{\text{in}} - F}{S} + 1 \quad (5.12)$$

$$H_{\text{out}} = \frac{H_{\text{in}} - F}{S} + 1 \quad (5.13)$$

At the end of the procedure the filter has computed a tensor resulting from this operation. The scheme of the convolutional operation is reported in Fig. 5.3a. Clearly, applying multiple filters results in obtaining multiple output tensors and we have something similar to 5.3b.




(a) Example of convolutional operation of a (3×3) filter K applied to a (7×7) input image I . The numbers inside the filter K are known as kernel's weights. The result is a (5×5) tensor

(b) Multiple features derived from the input. The smallest squares inside the white box represent the different filters, the gray box is the output of each filter.

Figure 5.3 Illustrations of a convolutional operation (a) and the result of a series of kernel applied to an input image (b).

The Convolutional Neural Network architecture consists in piling up several filter blocks, obtaining what is known as a convolutional block. Combining several convolutional blocks, along with other layers with different functionalities that are needed to have a properly working-CNN, a feature map is created. An example of this is reported in Fig. 5.4a. The convolutional operation ameliorate significantly the capabilities of such network in imaging recognition.



(a) A sequence of convolutional blocks (grey)—each comprising convolution, optional downsampling, and a non-linear activation—progressively extracts higher-level features while reducing spatial resolution. The final feature maps are flattened and passed to fully connected layers (colored block, right) that perform classification.

(b) Residual Block - shortcut connection. The input x is summed to the result of the inner layer of the Residual block F .

Figure 5.4 Illustrations of (a) the Convolutional Neural Network structure and (b) the shortcut connection that characterize the Residual Block

Similar to what we have seen for the NN case, a CNN learns its kernel weights via backpropagation. One drawback of such architecture is the gradient vanishing problem. In Eq. 5.7 we see how E is related to the weights w_{ji} , and the training procedure is done by minimizing the error function by

selecting the best set of weights. For really deep network, such as those used today for these tasks, the computation of gradient can be prone to vanishing, since in each passage the shrink of this quantity can be dramatic, spanning several orders of magnitude between the first and last layers. This lead to unstable or even to non-converging of training. To overcome this problem another architecture has been implemented: the Residual Block [164]. We saw that the Convolutional block consists in a set of filters that map the input tensor into an output tensor, in general then we want that a mapping function $\mathbf{H}(\mathbf{x})$ has to be learned, since it represents any transformation the network needs to perform at each passage. The Residual Blocks are based on the idea that for the network is easier to learn the residual function $\mathbf{F}(\mathbf{x})$ than $\mathbf{H}(\mathbf{x})$, where the residual function given the block's input \mathbf{x} is defined as

$$\mathbf{F}(\mathbf{x}) = \mathbf{H}(\mathbf{x}) - \mathbf{x} \quad (5.14)$$

Using this procedure, the network is no more prone to gradient vanishing problem and deep network capable in learning in very efficient way can be implemented. This procedure is implemented using the shortcut connection, that means to connect the input directly to the processed output of the convolutional blocks, depicted in 5.4b. This technique represent the state-of-the art for image recognition problems, and it is successfully implemented for architecture such are ResNet-152, where 152 layers have been stacked together. In this work we will use this model for gravitational-wave detection.

5.2 The Time-Template map

In the previous section we briefly introduced the most-used technique for image recognition problems. Here we describe how we can apply this tool for gravitational-wave detection. To do so, we need to figure out which kind of pictures will be created to train and test the Residual Network. In Chapter 2, we defined the matched filtering as the observable that quantifies the loudness of a signal that is present in the data strain. In Chapter 3 we saw also that the pipelines that run in the data acquisition campaign and provide the list of triggers, require template banks that span the possible parameters' space of the signal. The template bank must cover well enough the parameter space since any difference between the template we are matching the data and the signal will have an impact of the SNR. The mismatch quantify this loss in SNR due to the difference between waveforms. Usually, the maximum mismatch permitted is 0.3%. Recalling the definition given, if we consider the data strain provided by the detectors as $s(t)$ and the template to match against the data as $h(t)$, we have the time series defined as

$$\rho(t) = 4 \left| \int_0^\infty \frac{\tilde{h}^*(f)\tilde{s}(t)}{S_n(f)} e^{2i\pi ft} df \right| \quad (5.15)$$

This quantity will show a maximum in correspondence of the signal, as depicted in Fig. 2.6. The classical searches match-filter the data against the templates in the template bank, and look after the maximum SNR value computing the $\rho(t)$ for each template.

As stated above, any mismatch between the template and the signal in the data will lead to a loss in SNR. But this is not the solely reason to require efficient banks in covering the parameters space. In Chapter 2, we introduced the χ^2 rejection test defined in [78]. This test is fundamental for the non-Gaussian noise behavior of the interferometers. The test consists constructing a statistic which indicates if the filter output is consistent with the expected signal. This is done by computing the χ^2 distribution over a specified number of frequency intervals N_b , defined such that the expected signal contributions in each frequency bin are equal. In our study, we use $N_b = 4$. The number of bins is a function of the template (see [165] for reference), here we fix this number to 4 for practical reasons, even though this choice is not the optimal one.

$$\chi_r^2 = \frac{\chi^2}{(2N_b) - 2}. \quad (5.16)$$

The matched-filtering SNR $\rho(t)$ is then re-weighted using this χ_r^2 value according to this expression

$$\rho(t)_{\text{rw}} = \frac{\rho(t)}{\left[\frac{(1+\chi_r^2)^2}{2}\right]^{\frac{1}{6}}} \quad (5.17)$$

this new statistics incorporate the SNR information corrected with the χ^2 -test. The impact of this quantity in the ranking-statistics is quite evident: in Fig. 5.5 we consider a template h_0 for a non-spinning BBH with masses h_0 with $m_1 = m_2 = 15M_\odot$. Then, in Fig. 5.5a, we compute the SNR time series $\rho(t)$ for a template that perfectly matches the injected one. We see that the χ^2 -test has a peak in $\rho_{\text{rw}}(t)$ in correspondence to the maximum of ρ . If instead a glitch is present in the data, the behavior will be completely different as shown in Fig. 5.5b, where we injected a sine-Gaussian glitch in the strain. The χ^2 -test shows an impactful effect on glitch suppression in the ranking-statistics ρ_{rw} . What must be taken into account in any case is that the template must match very well the signal, otherwise the χ^2 -test will push down the ranking-statistics, possibly invalidating the detection even if a signal is present in the data. For instance, in Fig. 5.5c we matched the same data from Fig. 5.5a that contains a signal, but considered a slightly different template h_1 with $m_1^* = 15.2, m_2^* = 14.8M_\odot$ and $\chi_1^2 = \chi_2^2 = 0.02$. Even though the match between h_0 and h_1 is 0.97, we see that the χ^2 test dramatically impacts the detection reducing the ranking statistic by half. It follows that the choice of the template bank is a fundamental step in CBC searches, as we saw in Chapter 3 both to well recover the signal in the data and avoiding an extreme loss of detection efficiency due to the χ^2 test.




(a) Perfect template–signal match. $\rho_{\max} = 36.9$, $\rho_{\text{rw},\max} = 36.1$.(b) Template matched to a sine–Gaussian glitch. $\rho_{\max} = 15.1$, $\rho_{\text{rw},\max} = 5.1$.(c) Template–signal mismatch. $\rho_{\max} = 36.2$, $\rho_{\text{rw},\max} = 15.5$.

Figure 5.5 Effect of the χ^2 test on $\rho(t)$ for different template–data configurations. (a) Perfect match, (b) template matched to a sine–Gaussian glitch, and (c) template–signal mismatch. The reweighted statistic ρ_{rw} penalizes non–signal–like transients, significantly reducing their significance.

The χ^2 -test must be used carefully and requires a bank covering well-enough that the mismatch does not impact severely the ranking-statistics. In this work we ask ourselves if it is possible to perform a GW search with a reduced number of templates and avoiding the χ^2 -test in a parameter space that is compatible with the actual searches.

With this goal in mind, here we construct a one-dimensional template bank targeting equal-mass binary system ($m_1 = m_2$), including both BNS and BBH. The parameter space is defined in term of chirp-mass M_c .

At leading (OPN/Newtonian) order in the stationary-phase approximation, the inspiral phase of the signal depends only on M_c . After maximizing over coalescence time and phase, templates sharing the same M_c yield nearly identical matched-filter SNR even for systems with different (m_1, m_2) , which motivates a one-dimensional bank in M_c . Deviations due to the symmetric mass ratio $\eta = \frac{m_1 m_2}{(m_1 + m_2)^2}$ enter at higher PN orders (e.g. 1PN and above). The GW parameter space is continuous, whereas the template bank must be constructed from a discrete set of points. This discretization introduces a mismatch, as the templates in the bank typically do not match the signal exactly. This resulting loss in SNR is quantified by the Fitting Factor (FF), which measures the fractional loss in SNR due to the bank's sparsity. Here we choose a minimum FF loss of 0.3%. The template placement algorithm determines the optimal chirp-mass spacing between adjacent templates by ensuring that the match between them is equal to the minimum required match, here set to 0.97. The match is computed using the detector power spectral density (PSD), here modeled by aLIGOZeroDetHighPower that is an ideal design sensitivity of Advanced LIGO [8], which correspond to a BNS range of 173Mpc¹. To compute the match, the waveforms are generated using approximants obtained from solving the Einstein field equations for binary systems in General Relativity. In this work the approximant used is SEOBNRv5_ROM [166]. The SEOBNR framework combines post-Newtonian theory, black-hole perturbation theory, and numerical-relativity calibration within a unified formalism to model the full inspiral–merger–ringdown evolution of binary coalescences. The v5 generation incorporates improved conservative dynamics and radiation reaction, resulting in increased accuracy across a wide region of the parameter space for non-precessing, spinning binary black holes. The reduced-order implementation significantly accelerates waveform generation while preserving the fidelity of the underlying model, making it particularly well suited for match computations and large-scale parameter studies [167, 100]. Applying this procedure results in a one-dimensional template bank consisting in 5,442 templates. The masses span in the interval $m \in [1 - 50]M_\odot$, with a corresponding chirp-mass interval $M_c = [0.87 - 43.5]M_\odot$.

The matched filtering process produces an SNR time series for each template in the template bank described above, with increasing values of $\rho(t)$ in correspondence to the presence of an injected signal. In this work, we explore the possibility of identifying a candidate event by simultaneously analyzing the $\rho(t)$ values across the entire template bank within a time window of few seconds.


¹This range consists in the distance (in Mpc) where a BNS signal of $1.4 - 1.4M_\odot$ masses can be detected with SNR = 8.

We expect correlations among templates — induced by matched filtering — to produce structured patterns across the bank that encode both the presence of a signal and its parameters. To clarify this, let us consider a non-spinning BNS signal with component masses $m_1 = m_2 = 2.5M_\odot$ (chirp mass $M_c^{\text{ref}} = 2.17M_\odot$) merging at $t_0 = 40\text{s}$. In time-frequency representations, templates with lower chirp mass than the injection trace a slower inspiral and therefore their chirp tracks intersect the reference track after t_0 , see Fig. 5.6a; conversely, higher chirp-mass templates chirp faster and intersect the reference track before t_0 , see Fig. 5.6b. This ordering reflects how M_c controls the alignment of tracks in the time-frequency plane. The same geometry appears in the SNR time series: higher values in $\rho(t)$ from lower- M_c templates is shifted to times after t_0 , whereas higher- M_c templates peak before t_0 .


Next, we construct what we define as Time-Template SNR map (TT-SNR Map, hereafter TT-Map). Each template of the bank is matched-filtered against the selected stretch of data and each template yields a different SNR time series $\rho(t)$, as previously discussed. By stacking the $\rho(t)$ time series from all templates, we obtain the TT-Map, which represents the evolution of the SNR across both time and template dimensions. If a signal is present in the data, then the TT-Map exhibits characteristic structures that encode its presence and morphology. In addition, the presence of glitch will produce completely different correlation between templates, something that has already been studied in [168], but here will be exploited as an application for signal detection. Such features can be automatically recognized and classified by machine learning algorithm, such as a Convolutional Neural Network (CNN). Examples of these TT-Maps are shown in Fig. 5.7

In this study, to construct the TT-SNR Maps, we inject a GW signal into 80-second-long data segments using SEOBNRv4_opt approximant in Gaussian noise colored by aLIGOZeroDetHighPower power spectral density. These signals are then matched-filtered against the one-dimensional template bank generated with SEOBNRv4_ROM model, using a frequency range $\Delta f \in [40 - 2048]\text{Hz}$ with the same PSD used to color the Gaussian noise and construct the bank. To fasten up the computation, the data are downsampled from 16384 to 2048Hz. Once the $\rho(t)$ time series is computed for each template we extract a 4-second time window centered on the injection time. This produces an $N \times M$ matrix, where N is the number of templates in the bank while M is the number of time samples in the window, i.e. 4×2048 . The elements of the matrix are the $\rho(t)$ values. Since CNN training is sensitive to input image dimension, we compress the TT-SNR-Map into a 512×256 greyscale image. To standardize the representation, we map the SNR time series in greyscale using a logarithmic normalization with bounds $[b_{\min} = 0.1 - b_{\max} = 100]$, that is $\bar{\rho}(t) = \frac{\log(\rho(t)) - \log(b_{\min})}{\log(b_{\max}) - \log(b_{\min})}$, and then we clip $\bar{\rho}(t)$ to the unit interval. In the resulting greyscale image, high-SNR values appears as brighter pixels, while low-SNR corresponds to darker shades of gray.

The core idea is that the CNN can learn the features-rich representation provided by the TT-Map to distinguish between noise and injected signals. This reframes the classical GW triggering problem as an image recognition task, where CNNs are known to perform well. Each image is labeled as 'noise' or 'injection' depending on whether it contains purely Gaussian noise or an injected signal. The CNN is trained to output a score $r \in [0, 1]$ — also identified as CNN-output — with values near 0 indicating a noise-like classification and values near 1 indicating the presence of an injection. For



(a) Spectrogram of systems with $M_c = 0.95M_\odot$ merging time at $t_1 = 38s$ (red line), $M_c^{\text{ref}} = 2.17M_\odot$ merging time at $t_0 = 40s$ (green line) and $M_c = 0.95M_\odot$ merging time at $t_2 = 42s$ (blue line).



(b) Spectrogram of systems with $M_c = 4.35M_\odot$ merging time at $t_1 = 38s$ (red line), $M_c^{\text{ref}} = 2.17M_\odot$ merging time at $t_0 = 40s$ (green line) and $M_c = 4.35M_\odot$ merging time at $t_2 = 42s$ (blue line).

Figure 5.6 Spectrograms showing template–template correlations. We used a reference BNS with $M_c^{\text{ref}} = 2.17M_\odot$ and compare systems with chirp-mass below (top) and above (bottom) the reference. Each panel shows the time–frequency representation of the reference signal (central track) together with two realizations of the same template chirp mass, shifted in time with respect to the reference coalescence time t_0 . The presence of three chirp tracks therefore reflects the reference signal and two time-shifted instances of the template waveform, used to illustrate how the relative merger time determines whether the template intersects the reference chirp. The reference chirp track is intersected by a lower M_c system only if it merges after the reference time t_0 ; conversely, a higher M_c system crosses the reference track if it merges before t_0 .




Figure 5.7 Examples of the TT-SNR-Map obtained by piling up the $\rho(t)$. On Top the picture corresponding to Gaussian Noise solely, on the center an injected signal and on the bottom an injected signal with a glitch superimposed. The horizontal and vertical axis represent, respectively, the time interval around the injection time and the chirp-mass of the template used to matched-filter the data.

training, 70% of the dataset is used, while the remaining 30% is equally split between validation and test sets.

5.3 ResNet, training and results

In this work, we adopt a specific architecture known as Residual Network (ResNet) [164]. As described at the beginning of the chapter, this model is explicitly designed for image recognition tasks and has been optimized for such purposes through the Residual Block architecture, that is the state-of-the-art algorithm for image recognition. In our case, this flexible model is trained specifically to distinguish the TT-Map generated from noise and those containing signal. The structure of the network is depicted in Fig. 5.8a.

The input consists of a 512×256 greyscale image. It first passes through a convolutional layer with 8 filters of kernel size (3×3) and a stride of 2 (i.e. the step size of the convolution). This produces a $256 \times 128 \times 8$ output tensor, which is then normalized using Batch Normalization (BN). The batch normalization takes as input the features map F , of $(H \times W)$ dimension with coordinates (h, w) , and normalizes it: in our example the BN computing the mean μ and variance σ gives the normalized feature map $\hat{F}(h, w)$

$$\mu = \frac{1}{H \cdot W} \sum_h^H \sum_w^W F_{h,w} \quad (5.18)$$

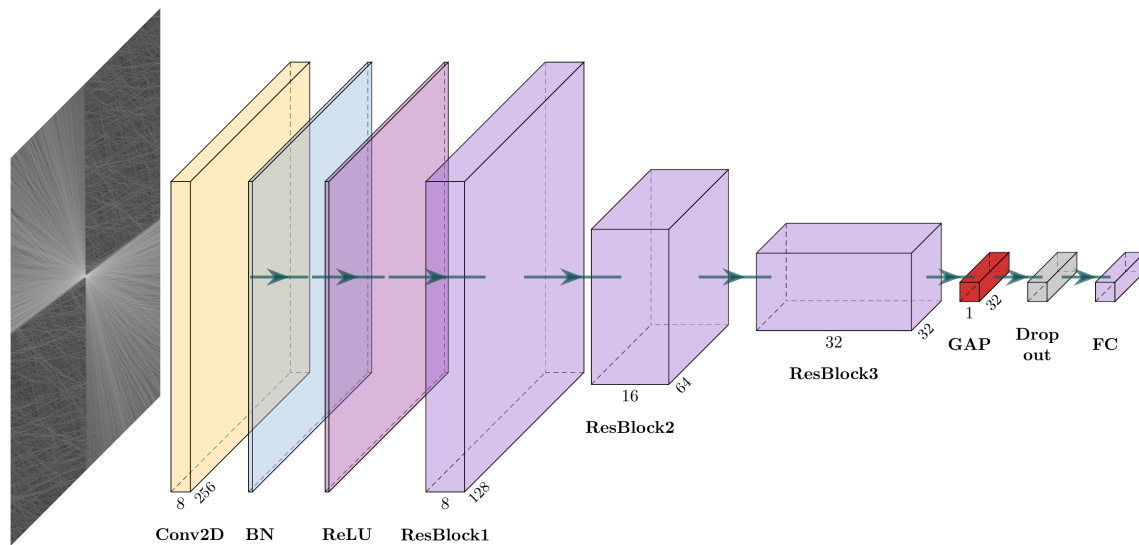
$$\sigma = \frac{1}{H \cdot W} \sum_h^H \sum_w^W (F_{h,w} - \mu)^2 \quad (5.19)$$

$$\hat{F}(h, w) = \frac{F_{h,w} - \mu}{\sigma} \quad (5.20)$$


The data pass through three sequential Residual Blocks, schematized in Fig. 5.8b, that progressively extract higher-level features. After these blocks, a Global Average Pooling (GAP) layer aggregates the most relevant informations by averaging each feature map into a single value, producing a feature vector following this equation

$$P = \frac{1}{MN} \sum_i^{M-1} \sum_j^{N-1} F(i, j) \quad (5.21)$$

Where F is the feature map in input of $(M \times N)$ dimension. This output is then activated with a Rectifier Linear Unit (ReLU), an activation function similar to Eq. 5.2. After these three Residual Blocks, this vector is then passed to a fully connected (FC) layer that outputs a classification score for the input image. To mitigate overfitting, a dropout layer with a dropout rate of $p = 0.3$ is applied between GAP and FC layers, randomly setting the chosen fraction of the input to zero during the training, this procedure is also known as regularization. The network is relatively lightweight, consisting of approximately 20,000 trainable parameters and it is optimized using Adam algorithm for 20 epochs. Adam adaptively adjusts learning rates based on estimates of the first and second moments of the gradients of the loss function.



(a) EasyResNet architecture. A convolutional stem (Conv2D-BN-ReLU) processes the input and feeds three residual blocks (ResBlock1-3). The final feature maps are reduced by global average pooling (GAP) to a feature vector, regularized with dropout, and passed to a fully connected (FC) layer for classification. Arrows indicate data flow; numbers below the blocks denote channel counts (and, where shown, spatial dimensions).



(b) Residual block. The lower arrow represents the input \mathbf{x} , producing $\mathbf{F}(\mathbf{x})$ passing through a sequence of convolutions and normalizations. The upper one carries the identity \mathbf{x} . The green \oplus denotes element-wise addition, yielding the block output $\mathbf{y} = \mathbf{F}(\mathbf{x}) + \mathbf{x}$. Arrow directions indicate data flow through the layers (Conv2D-BN-ReLU-Conv2D-BN-Conv2D in this example).

Figure 5.8 Architecture of the Residual Network used in this work (Top) and focus on the Residual Block (Bottom).



5.3.1 Simulation 1

As first dataset, we considered the idealized case of purely Gaussian noise with injected BNS signals. This dataset contains in 15,000 images for each class, 'noise' and 'injection', and the train-validation-test split follows the procedure described above (70-15-15). The injected signals were generated by sampling the component masses uniformly (\mathcal{U}) in the interval $m_{1,2} \in \mathcal{U}[1.4 - 3]M_{\odot}$. No spins were assigned to the system components. The distance distribution of the sources follows a power-law probability density distribution $\sim d^2$ with distance drawn from the range $d \in [50 - 1000]$ Mpc. The lower bound at 50Mpc is chosen to avoid a overabundance of extremely high SNR events, which are trivially detectable and provide limited information for training the network. The upper bound 1000Mpc extends well beyond the typical detection horizon for BNS systems, ensuring the inclusion of marginal and sub-threshold signals. This wide distance range allows the dataset to span smoothly from clearly detectable events to effectively undetectable signals, encouraging the network to learn a continuous mapping between signal morphology and detectability. Although distances close to 1000 Mpc are not astrophysically representative for observed BNS mergers, their inclusion is motivated by methodological considerations, namely to expose the network to a realistic mixture of loud, weak, and noise-dominated signals consistent with a volumetric ($\propto d^2$) distribution. The metric used to assess the performance of the network is the Cumulative Density Function (CDF) for the population t — noise (t=n) or injection (t=s) — defined as

$$\text{CDF}_t = \frac{1}{N_t} \sum_i^{N_t} \theta(r_i - \hat{r}) \quad (5.22)$$

Here, N_t is the number of samples in the test dataset for the population t , θ denotes the Heaviside step function, and r_i is the score (or CNN-output) assigned by the network to the i -th sample. This score r_i is compared to a varying threshold \hat{r} ; the CDF thus quantifies the fraction of events in population t with scores above the threshold. In other words, it characterizes the score distribution of each population as function of r . A well-trained network is expected to assign scores $r \sim 0$ to noise-labeled images and $r \sim 1$ to injection-labeled TT-map.

Figure 5.9 shows the results of the test dataset classification performed by the EasyResNet. On Fig. 5.9a we see the correlation between the CNN-output, encoded by the colorbar, the distance and the optimal SNR ρ_{opt} . We observe that the injected signals follow the power-law density distribution: few nearby sources and many distant ones. Nearby injections have higher SNR, and their TT-Maps will manifestly show the structures associated to the injections. The farther we go, the fader this map will be, since fainter will be also the SNR, see for reference Fig. 5.10. In this way we see that the Residual Network assigns high score to loud (near) signals, while it fails in recognizing if they are too far away. In Fig. 5.9b instead we see the CDF distribution: the network fails to separate the two populations, and for lower values ($r < 0.3$) the populations completely overlap. This degeneracy is expected, as distant injections — or equivalently, signals with low optimal SNR ρ_{opt} — cannot be reliably distinguished from the noise distribution.

(a) Distribution of injected signals for ρ_{opt} - distance Mpc - CNN output.

(b) CDFs for Noise - Signal population

Figure 5.9 Results for purely Gaussian Noise and BNS injections.




Figure 5.10 BNS signals TT-Map injected, respectively, at 50 (a), 100 (b), 200 (c) and 400 (d) Mpc. We see that the farther the signal is, the fader the TT-Map will result, consequently really far away signals cannot be distinguished from pure Gaussian noise.

To evaluate and compare the classification performance of the CNN and the traditional ranking statistics, we use the Receiver Operating Characteristic curve (ROC). This curve is obtained by considering the False Alarm Probability FAP and the Number of detection N_d as function of a varying classification threshold. The FAP and N_d are defined as

$$\text{FAP} = \frac{1}{N_n} \sum_i^{N_n} \theta(k_i - \hat{k}) \quad (5.23)$$

$$N_d = \sum_i^{N_s} \theta(k_i - \hat{k}) \quad (5.24)$$




Figure 5.11 Simulation 1: ROC curves for r (CNN-output) and ρ^{\max} statistics.

Here N_n (N_s) are the total number of noise (signal) in the test dataset. The variable k denotes the statistics under consideration: either the matched-filter maximum ρ^{\max} , that is the maximum value of ρ in the time-window considered to build the TT-Map, or the CNN-output score r . The threshold \hat{k} is varied to generate the ROC curve, which expresses the number of detection N_d as function of the FAP. In Fig. 5.11 is depicted the ROC curve obtained for this first simulation.

We see that the ρ curve is always above the statistics provided by the ResNet for classifying the TT-Map. This means that at give rate of false positive, the SNR is a better statistics for triggering potential gravitational-wave candidates. This result is kind of expected since the matched-filter is an optimal filter in pure Gaussian noise, that is exactly this case, so even a well trained Network find difficulties in out-performing the ρ statistics. In addition, our ResNet is trained on a downsampled representation: the TT-Map inputs are 512×256 pixels, whereas ρ_{\max} is computed by scanning the full-resolution time series over the complete template grid and taking the maximum across all samples. This dimensionality reduction discards information that the matched-filter search exploits, putting the CNN at a disadvantage.

5.3.2 Simulation 2

The second simulation considered in this work includes a dataset contaminated by non-Gaussian noise, which affects both the 'noise' and 'injections' classes. The dataset consists in 20,000 TT-Maps per class, divided in training-validation and tested as previously described. In this case, the injected

signals include both BNS and BBH systems. The component masses for BNS and BBH injections were drawn from two different distributions. For the BNS systems, we retained a uniform distribution for the component masses as before. For BBH systems, the primary mass m_1 was sampled from an astrophysically motivated distribution [150], modeled as a power law with slope $\alpha = 3.5$ and a Gaussian peak centered at $34M_\odot$. The secondary mass m_2 was computed via the mass ratio $q = m_2/m_1$, where $q \sim \mathcal{U}[0.15, 1]$ and the final sampling followed the transformation $q = q^{0.5}$ to favor asymmetric systems. The allowed range for the primary mass was $m_1 \in [3 - 200]M_\odot$. Both BNS and BBH signals amplitude were scaled according to the chirp-distance, defined as

$$d_{\text{chirp}} = \left(\frac{1.22}{M_c} \right)^{\frac{5}{6}} \cdot d \quad (5.25)$$

with d distributed according to a second-order power law, as in the previous case. This scaling accounts for the fact that the observable amplitude of a gravitational-wave signal depends on the chirp-mass M_c , making d_{chirp} a more appropriate parameter for distance-scaling when comparing signals of different masses. In this simulation, we injected sine-Gaussian glitches randomly into 30% of data segments to emulate the non-Gaussian transient noise. A sine-Gaussian in time domain can be written as

$$g(t) = Ae^{\frac{(t-t_0)^2}{\tau^2}} \cos(2\pi f_0 t + \phi_0) \quad (5.26)$$

In this expression, A identifies the amplitude and $\tau = \frac{Q}{2\pi f_0}$ the duration of the pulse, where Q is the dimensionless quality factor. The parameters f_0 and t_0 specify the central frequency and central time of the sine-Gaussian, respectively, while ϕ_0 is the phase at $t = 0$. Since it is known that the matched-filtering response to such transients can extend significantly beyond the glitch time itself, we injected the sine-Gaussian glitches uniformly in a time window $\Delta t \in [-63, 0.5]$ s, taking the nominal injection time of the binary system at $t = 0$. This simulates the possible effect of the glitches even when they occur well before the actual signal. The remaining sine-Gaussian parameters were sampled from uniform distributions \mathcal{U} , with the corresponding ranges listed in Table 5.1. An example of injected signal with sine-Gaussian glitch is depicted in Fig. 5.12

Table 5.1 Parameter ranges used for the injected sine-Gaussian signals.

Parameter	Value range
Amplitude A	$[1-9] \times 10^{-22}$
Central frequency f_0	$[20, 2048]$ Hz
Quality factor Q	$[3, 400]$
Initial phase ϕ_0	$[0, 6]$

In this section we construct the TT-Map as done in Simulation 1, and compare the CNN-output score r with a ranking statistic that incorporates the χ^2 -test to suppress glitch-induced SNR, the $\rho_{\text{rw}}(t)$. Once the matched-filtering procedure is complete and the ranking-statistics is computed, we




Figure 5.12 Injected BNS signal in Gaussian noise with sine-Gaussian glitch (top) and zoom over glitch (bottom)

then identify the the maximum value $\rho_{\text{rw}}^{\text{max}}$ within the 4-seconds time window and compare it to the CNN-output r for the corresponding TT-Map.

We expect that the reweighted statistics $\rho_{\text{rw}}^{\text{max}}$, which incorporates the χ_r^2 test, will be significantly more effective at separating noise from signal than the raw matched-filter maximum ρ^{max} . We then compare the ROC curves obtained using $\rho_{\text{rw}}^{\text{max}}$ and the CNN-output r statistics to quantify the network's ability to distinguish between glitches and true signals relative to the classical approach. Results are reported in Fig. 5.13

In the results we observe that the CNN-output outperforms the raw matched-filter statistic ρ statistics, particularly for low false alarm probabilities. However, the reweighted statistic $\rho_{\text{rw}}^{\text{max}}$ remains the most effective overall.

We would like to spend a couple of words regarding the $\chi_r^2(t)$ effect to the TT-Map. It is instructive in fact to see how this observable modulates the $\rho(t)$ in building the $\rho_{\text{rw}}(t)$ for each template in the TT-Map (Fig. 5.14). Here we observe an effect similar to what we have seen in Fig. 5.5 but instead of a single template a collective bank behavior is represented.

As an illustrative case, we inject a sine–Gaussian glitch near the target signal. Fig. 5.14a shows the TT-Map constructed from $\rho(t)$: the signal exhibits the characteristic 'butterfly' pattern in the SNR time series, whereas the glitch has a distinct morphology. A zoom around the merger time reveal a cluster of large $\rho(t)$ values. In Fig. 5.14b the corresponding χ_r^2 -test is reported and the zoom shows high values across most templates, but low for the most similar template, visible as a darker pixel at the center of the zoomed region. Finally, Fig. 5.14c displays the TT-Map built from ρ_{rw} , where the




Figure 5.13 Simulation 2: ROC curves for ρ_{rw}^{\max} , r (CNN-output) and ρ^{\max} statistics.

ranking-statistics suppresses high χ_r^2 (glitch-like) templates and concentrates the peak on the most similar template².

5.3.3 Simulation 3

In this final study, in addition to sineGaussian glitches present in 30% of data, we included also physical informations in the injected signals that are not represented in the template bank. Specifically, spin effects were incorporated by assigning aligned-spin components drawn from uniform distributions. For Neutron-stars, we used a conservative spin range of $\chi_{z,NS}^{1,2} \in [-0.5, 0.5]$ [169], while for Black-Holes we adopted $\chi_{z,BH}^{1,2} \in [-0.9, 0.9]$, reflecting the typically higher spin magnitudes of BH. The masses and the distances follow the distributions used for Simulation 2. As in the previous studies, the dataset was divided into training-validation-test in a 70 – 15 – 15 split. However, this time the total dataset size was significantly larger, consisting in 40,000 TT-Maps per class for both noise and injection samples. The relatively smallness of the network allowed a training time around ~ 140 minutes. The performance comparison followed the same procedure as in the previous section, and is represented by the ROC curves in Fig. 5.15.

In this case the physical effect of spin-absent from the template bank causes the χ^2 test to down-rank the SNR associated with the true injections. In contrast, EasyResNet is not affected by this

²Here the pictures are not downsampled, so we effectively see the most similar template in the time-interval used to compute the matched-filtering.

mismatch and achieves performance comparable to that of the matched-filtering search, even in presence of spin.

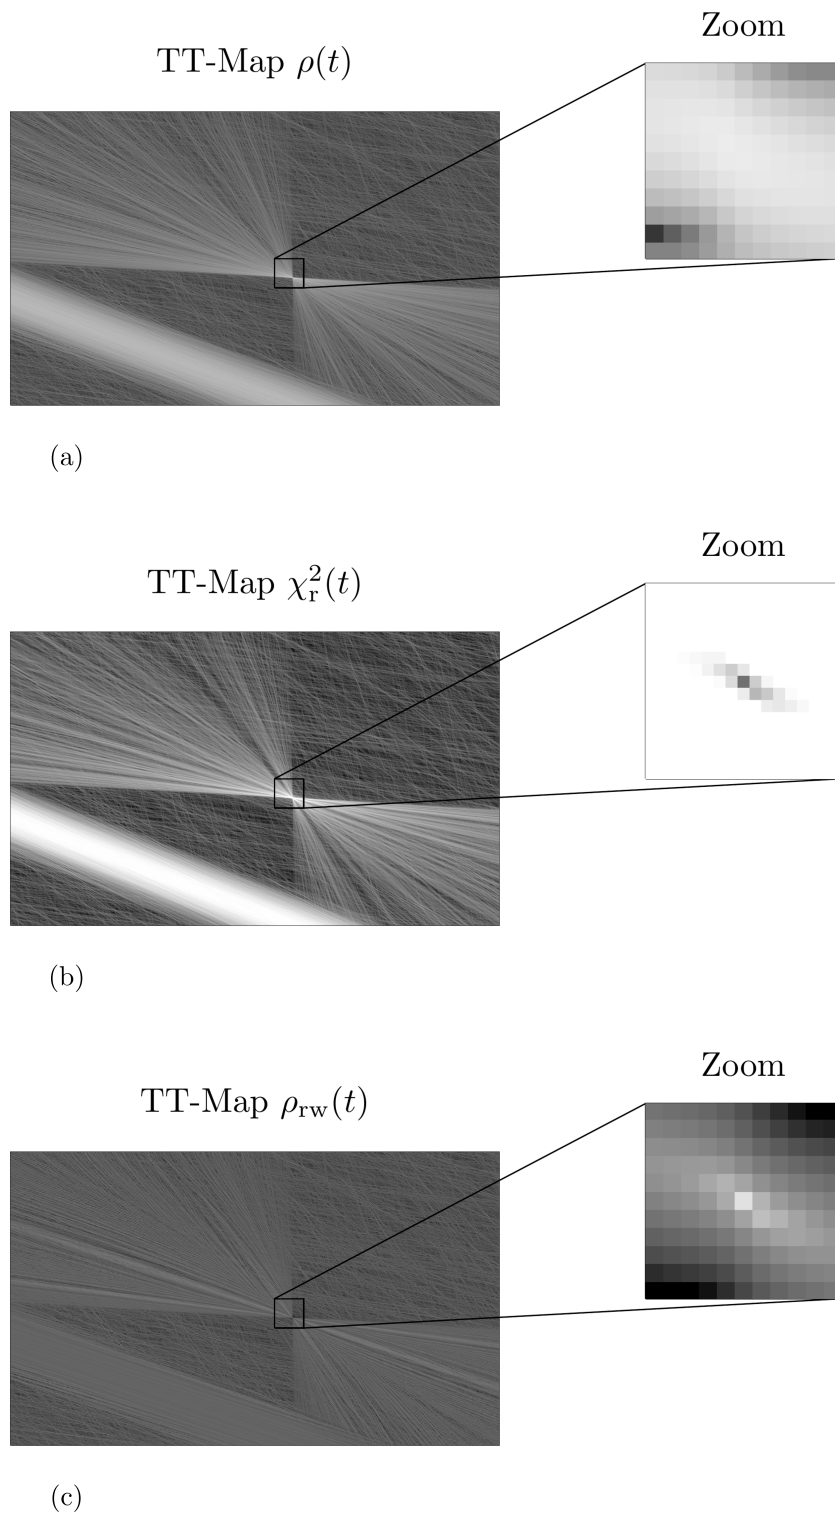


Figure 5.14 TT-Map constructed considering $\rho(t)$ (top), $\chi_r^2(t)$ (center) and $\rho_{rw}(t)$ (bottom) statistics. The horizontal and vertical axes are the same as in Fig. 5.7.




Figure 5.15 Simulation 3: ROC curves for $\rho_{\text{rw}}^{\text{max}}$, r (CNN-output) and ρ^{max} statistics.

Chapter 6

Conclusions

Gravitational-wave astronomy is a rapidly expanding and highly promising research field. Since the first detection of GWs, with the observation of GW150914, more than 218 events have been observed up to the last observation campaign O4. The great scientific advancement achieved in this domain is remarkable: the theory of General Relativity, formulated by Einstein over a century ago, continues to demonstrate its robustness, while discoveries concerning to the nature of astrophysical objects — such as Black-Holes or Neutron Stars — are deeply linked to the gravitational-wave observations and detections. In this context, the compact binary coalescence searches performed by pipelines such as MBTA, pyCBC and GstLAL play a crucial role. As a large fraction of the triggers reported by the LIGO-Virgo-Kagra collaboration originate from these algorithms, ensuring their optimal performance and scientific reliability is of paramount importance.

In Chapter 3, we explored a novel approach to construct the template banks required by MBTA for performing matched-filtering on interferometric data. Since this pipeline computes matches in multiple frequency bands, building the corresponding banks for each specific bandwidth is a fundamental step. By studying the impact of the frequency-cut f_{opt} on the construction of these banks, we implemented a MapReduce-based algorithm to generate them, achieving effectively covering banks in only a fraction of the time needed by the standard approach. The main drawback of this method is that the resulting number of templates are approximately twice time higher than that obtained with the classical technique. Further improvements in this direction could be achieved by adopting alternative strategies — for instance, by placing an optimal template in considered bandwidth and associating with it the templates within a certain interval of parameter space, such as the chirp-mass M_c . Experiments following this line of investigation are currently being carried, and promising preliminary results are emerging.

In addition to template bank construction, this thesis also explores the application machine learning techniques to compact binary coalescence searches. Although this field has been investigated in previous studies — such as the classification of single triggers using Random Forests for pyCBC and GstLAL — in Chapter 4 we extended the applicability of such methods for the MBTA pipeline. Specifically, we analyzed single triggers from O4a data acquisition campaign and coincidence triggers

of O3a and O3b. In both cases, we observed an improvement in discrimination between signal and noise triggers when comparing the ranking-statistics obtained from MBTA with those derived from the Random Forests. In particular, the classification of single detection triggers shows a consistent enhancement across the entire parameter space, owing the inclusion of multiple features that capture both statistical and physical informations. For the double coincidence, the ranking-statistics were found to be consistent with those produced by the Random Forests, confirming that enforcing coincidence and coherence between triggers is already an effective approach to suppress transient noise. Furthermore, by exploiting non-parametric methods for probability density estimation, we computed the astrophysical probability p_{astro} based on the Random Forests statistics and validated its reliability against the events reported in the catalogues GWTC-2.1 and GWTC-3. The results demonstrate consistency between the astrophysical probabilities derived from the MBTA ranking-statistics and those obtained via the Random Forests classifier. Finally, an agnostic search among the all triggers was conducted using the trained model, resulting in the identification of a candidate with $p_{\text{astro}} = 0.92$ at GPS time 1240423628.7. This trigger is reported also in [154] with $p_{\text{astro}} = 0.41$ as GW190427_180650.

In Chapter 5, we conducted a proof-of-concept study exploring the application of deep learning to Gravitational Waves detection by exploiting signal-to-noise ratio time-series. Although the use of deep learning for gravitational-wave searches is not new — and the analysis of signal-to-noise ratio time-series has been investigated in previous works — our study introduced a novel, feature-rich representation of such time series designed to be learned by a Residual Network in a innovative way. To this end, we constructed an ad hoc template bank and developed the necessary preprocessing pipeline to transform the input data into a meaningful feature map. This representation effectively captures the correlation between templates and signal-to-noise ratio time series, enabling the network to learn the relevant signal characteristic. The resulting model demonstrated performance compatible that of the standard ranking-statistics method, with the advantage of not relying on the χ^2 -test for noise-glitch rejection, as the morphology of the input feature maps inherently encodes such informations. Extensive simulation campaigns were carried out to assess the model’s robustness. We first considered injected signals in Gaussian Noise, then introduced sine-Gaussian glitches, and finally tested injections containing physical effects not encoded in the template bank used for building the input pictures. These experiments confirmed the flexibility and generalization capability of the Residual Network in handling diverse signal and noise conditions.

Gravitational-wave astronomy is an exciting and rapidly evolving research field, with a promising future in both scientific discovery and technological development. The study of compact binary coalescences plays a central role in this research field, not only for its capability to detect and characterize gravitational-wave signals, but also as a platform for testing and applying innovative computational technologies to complex-datasets. In this spirit, this thesis has explored the application of machine learning methods to this challenging and stimulating area of research.

Acknowledgements

I would first like to thank my supervisor, Gianluca Maria Guidi, for its constant guidance and support throughout PhD course. Thank you for teaching me how to do research in gravitational-wave physics. Your suggestions and insights have been, are, and will continue to be fundamental for my scientific career.

Another person without whom this thesis would not be as rich as it is, is Tito dal Canton. He welcomed me to Paris and made me feel at home at IJCLab, providing an exceptionally stimulating and scientifically rewarding environment. Thank you for your clarity of thought and for the energy you brought to every discussion.

This work would not be possible without the support, discussions and suggestions provided by the MBTA team. In particular Thomas Sainrat and Florian Aubin for their constant and helpful presence.

I would also like to thank the Firenze–Urbino group, who supported me in my research and was always available for discussion and suggestions. Furthermore, I thank Kristian Toccacelo, for its valuable advices and scientific discussions.

Finally, I would like to thank my family and friends — from Perugia, Padova, Urbino and every other place in this strange world. Your presence has meant more than you know. And I thank Sara: although we sometimes spend long periods far apart, our bond is stronger than any distance, and your words are profoundly meaningful to me.

This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation, as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council

(NSTC) in Taiwan. Part of our simulations utilized the Virtual Data cloud computing system at IJCLab. We thank Michel Jouvin and Gerard Marchal-Duval for their prompt support and advice about this system.

Bibliography

- [1] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116(6):061102, 2016.
- [2] Campanelli M. et al. Accurate Evolutions of Orbiting Black-Hole Binaries without Excision. *Physical Review Letters*, 96, 2006.
- [3] Sxs:bbh:0305 waveform. <https://doi.org/10.5281/zenodo.10473621>, 2023. The waveform shown is SXS:BBH:0305, available for download at <https://doi.org/10.5281/zenodo.10473621>.
- [4] Mroué A. H. et al. Catalog of 174 Binary Black Hole Simulations for Gravitational Wave Astronomy. *Physical Review Letters*, 111, 2013.
- [5] Chatterji S. et al. Multiresolution techniques for the detection of gravitational-wave bursts. *Classical and Quantum Gravity*, 21(20), 2004.
- [6] California Institute of Technology LIGO Laboratory and MIT. Ligo lab | caltech. <https://www.ligo.caltech.edu>, 2025.
- [7] M. Bernardini et al. Active control hierarchy in Virgo superattenuator: The role of the inverted pendulum. In *2nd Edoardo Amaldi Conference on Gravitational Waves*, pages 334–338, 7 1997.
- [8] David Shoemaker. Advanced ligo anticipated sensitivity curves. Technical Report LIGO-T0900288-v3, LIGO Scientific Collaboration, 2010. Technical Report.
- [9] Trevor Hastie et al. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, New York, NY, 2 edition, 2009.
- [10] B. P. Abbott et al. GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119(16):161101, 2017.
- [11] R. Abbott et al. GW190521: A binary black hole merger with a total mass of $150 M_{\odot}$. *Phys. Rev. Lett.*, 125(10):101102, 2020.
- [12] B. P. Abbott et al. GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs. *Phys. Rev. X*, 9(3):031040, 2019.
- [13] R. Abbott et al. GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run. *Phys. Rev. X*, 11(2):021053, 2021.
- [14] R. Abbott et al. GWTC-2.1: Deep extended catalog of compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run. *Phys. Rev. D*, 109(2):022001, 2024.

- [15] R. Abbot et al. GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run. *Phys. Rev. X*, 13:041039, Dec 2023.
- [16] GWTC-4.0: Updating the Gravitational-Wave Transient Catalog with Observations from the First Part of the Fourth LIGO-Virgo-KAGRA Observing Run. 2025.
- [17] J.K Maxwell. A dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society of London*, 155:459–512, 1865. Paper read at a meeting of the Royal Society on 8 December 1864.
- [18] H. Poincaré. Sur la dynamique de l'électron. *Comptes Rendus des séances de l'Académie des Sciences*, 140:1504–1508, 1905.
- [19] A. Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (Berlin)*, pages 844–847, 1915.
- [20] Einstein A. and Rosen N. On gravitational waves. *Journal of the Franklin Institute*, 223(1):43–54, 1937.
- [21] M. Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press, 2007.
- [22] G. Pizzella. Search of gravitational waves with resonant detectors. *Journal of the Franklin Institute*, 9(9):81–96, 1989.
- [23] L. Cadonati et al. The AURIGA–LIGO joint burst search. *Classical and Quantum Gravity*, 22:1–11, 2005.
- [24] R. A. Hulse and J. H. Taylor. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195(L):L51–L53, 1975.
- [25] LIGO Scientific Collaboration et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, 2015.
- [26] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quantum Grav.*, 32(2):024001, 2015.
- [27] A. G. Abac et al. A search using GEO600 for gravitational waves coincident with fast radio bursts from SGR 1935+2154. *The Astrophysical Journal*, 977:255, 2024.
- [28] T. Akutsu et al. Overview of KAGRA: Detector design and construction history. *Progress of Theoretical and Experimental Physics*, page 05A101, 2021.
- [29] C. S. Unnikrishnan et al. LIGO-India: A decadal assessment on its scope, relevance, progress, and future. *International Journal of Modern Physics D*, 33(5 & 6):2450025, 2024.
- [30] A. Abac et al. The science of the Einstein Telescope. *arXiv preprint*, 2025.
- [31] E. D. Hall et al. Cosmic Explorer: A next-generation ground-based gravitational-wave observatory. *Galaxies*, 10(4):90, 2022.
- [32] M. Di Giovanni et al. Einstein telescope and cosmic explorer. *arXiv preprint*, 2025.
- [33] M. Colpi et al. LISA definition study report. *Experimental Astronomy*, 2025. in press.
- [34] T. W. Baumgarte and S. L. Shapiro. Numerical relativity and compact binaries. *Physics Reports*, 376(2):41–131, 2003.

- [35] A. Pound and B. Wardell. Black hole perturbation theory and gravitational self-force. *Springer Nature*, pages 1411–1529, 2022.
- [36] Kasen D. et al. Origin of the heavy elements in binary neutron-star mergers from a gravitational-wave event. *Nature*, 551(7678):80–84, October 2017.
- [37] B. P. Abbott et al. Gravitational waves and gamma-rays from a binary neutron star merger: Gw170817 and grb 170817a. *The Astrophysical Journal Letters*, 848(2):L13, October 2017.
- [38] Smartt S. J. and others. A kilonova as the electromagnetic counterpart to a gravitational-wave source. *Nature*, 551(7678):75–79, October 2017.
- [39] Radice D. et al. GW170817: Joint Constraint on the Neutron Star Equation of State from Multimessenger Observations. *The Astrophysical Journal Letters*, 852(L29), 2018.
- [40] Abbot B.P. et al. A gravitational-wave standard siren measurement of the Hubble constant. *Nature*, 551:85, 88, 2017.
- [41] Dall’Amico M. et al. GW190521 formation via three-body encounters in young massive star clusters. *MNRAS*, 508:3045–3054, 2021.
- [42] Morton S. L. et al. GW190521: A binary black hole merger inside an active galactic nucleus? *Physical Review D*, 108, 2023.
- [43] Abbott R. et al. Observation of gravitational waves from two neutron star-black hole coalescences. *The Astrophysical Journal Letters*, 915(1), 2021.
- [44] S Biscoveanu et al. Population properties and multimessenger prospects of neutron star–black hole mergers following GWTC-3. *MNRAS*, 518(4):5298–5312, 2022.
- [45] Colombo A. et al. Multi-messenger prospects for black hole – neutron star mergers in the O4 and O5 runs. *Astronomy and Astrophysics*, 686:A265, June 2024.
- [46] Abac A.G. et al. GW231123: A Binary Black Hole Merger with Total Mass $190 - 265M_{\odot}$. *The Astrophysical Journal Letters*, L25, 2025.
- [47] Paiella L. et al. Assembling GW231123 in star clusters through the combination of stellar binary evolution and hierarchical mergers. *The Astrophysical Journal Letters*, L54, 2025.
- [48] Gilles Theureau, S. Babak, Anne Berthureau, A. Chalumeau, S. Chen, Ismael Cognard, Mikel Falxa, Lucas Guillemot, and Antoine Petiteau. Pulsar timing arrays and gravitational waves : The first steps towards detection?
- [49] S. Valtolina et al. Regularizing the pulsar timing array likelihood: A path toward efficient analysis. *Phys. Rev. D*, 2025.
- [50] LISA Study Team. Lisa: Laser interferometer space antenna – mission summary. <https://sci.esa.int/web/lisa/>, accessed 2025.
- [51] M. Tinto et al. LISA sensitivities to gravitational waves from relativistic metric theories of gravity. *Phys. Rev. D*, 2010.
- [52] B. Sathyaprakash et al. Einstein gravitational-wave telescope: Conceptual design. In *Journal of Physics: Conference Series*, 2012.
- [53] E. D. Hall et al. Gravitational-wave physics with Cosmic Explorer: Limits to low frequencies. *Phys. Rev. D*, 2021.

- [54] E. D. Hall et al. Cosmic explorer: A next-generation ground-based gravitational-wave observatory. *Galaxies*, 2022.
- [55] Juhan F. et al. *Accretion Power in Astrophysics*. Cambridge University Press, Cambridge, UK, 3rd edition, 2002.
- [56] Alberto S. Gravitational wave emission from binary supermassive black holes. *Classical and Quantum Gravity*, 30(24):244009, 2013.
- [57] Maggiore M. et al. Science Case for the Einstein Telescope. *Journal of Cosmology and Astroparticle Physics*, 2020(03):050, 2020.
- [58] Marcoccia P. et al. Probing primordial black holes at high redshift with future gravitational wave detectors. *MNRSA*, 531, 2024.
- [59] Bromm V. Formation of the first stars. *Reports on Progress in Physics*, 76(11):112901, 2013.
- [60] Klessen R. et al. The first stars: Formation, properties, and impact. *Annual Review of Astronomy and Astrophysics*, 61:65–106, 2023.
- [61] A. A. Michelson and E. W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, 34(203):333–345, 1887.
- [62] Teviet Creighton. Tumbleweeds and airborne gravitational noise sources for ligo. *Classical and Quantum Gravity*, 25(12):125011, June 2008.
- [63] Peter R. Saulson. Thermal noise in mechanical experiments. *Physical Review D*, 42(8):2437–2445, 1990.
- [64] Gregory M. Harry et al. Thermal noise in interferometric gravitational wave detectors due to dielectric coatings. *Classical and Quantum Gravity*, 19(5):897–917, 2002.
- [65] Gregory M. Harry et al. Thermal noise from optical coatings in gravitational wave detectors. *Applied Optics*, 45(7):1569–1574, 2006.
- [66] Jessica Steinlechner. Development of mirror coatings for gravitational-wave detectors. *Philosophical Transactions of the Royal Society A*, 376(2120):20170282, 2018.
- [67] J. Franc. Mirror thermal noise in laser interferometer gravitational wave detectors operating at room and cryogenic temperature. Technical report, Einstein Telescope Project, 2009. ET-021-09.
- [68] Gabriele Vajente et al. Low mechanical loss coatings for reduced thermal noise in gravitational-wave detectors. *Physical Review Letters*, 127:071101, 2021.
- [69] R. X. Adhikari et al. Astrophysical science metrics for next-generation gravitational-wave detectors. *Classical and Quantum Gravity*, 36(14):245010, 2019.
- [70] M. Tinto. Antenna patterns of interferometric detectors of gravitational waves – II. elliptically and randomly polarized waves. *Monthly Notices of the Royal Astronomical Society*, 226(4):829–848, 1987.
- [71] B. J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Physical Review D*, 53(12):6749–6761, 1996.
- [72] T. Adams et al. Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era. *Classical and Quantum Gravity*, 33(17):175012, 2016.

- [73] F. Aubin et al. The MBTA pipeline for detecting compact binary coalescences in the third LIGO-Virgo observing run. *Classical and Quantum Gravity*, 38(09):095004, 2021.
- [74] C. All  n   et al. The MBTA pipeline for detecting compact binary coalescences in the fourth LIGO-Virgo-KAGRA observing run. *Classical and Quantum Gravity*, 42:105009, 2025.
- [75] The LIGO Scientific Collaboration et al. First low-latency LIGO+Virgo search for binary inspirals and their electromagnetic counterparts. *Astronomy and Astrophysics*, 541, 2012.
- [76] Sara Vallero et al. Gravitational wave alert generation infrastructure on your laptop. In *EPJ Web of Conferences*, volume 295, page 04022, 2024. CHEP 2023.
- [77] B. P. Abbott et al. A guide to ligo-virgo detector noise and extraction of transient gravitational-wave signals. *Classical and Quantum Gravity*, 37(5):055002, 2020.
- [78] B. Allen. χ^2 time-frequency discriminator for gravitational wave detection. *Physical Review D*, 71(062001):062001, 2005.
- [79] B. Abbott et al. Search for gravitational waves from binary inspirals in S3 and S4 LIGO data. *Physical Review D*, 77(6):062002, 2008.
- [80] C. Messick et al. Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data. *Physical Review D*, 95(4):042001, 2017.
- [81] A. Ray et al. When to point your telescopes: Gravitational wave trigger classification for real-time multi-messenger followup observations. *arXiv preprint*, 2023.
- [82] W. M. Farr et al. Counting and confusion: Bayesian rate estimation with multiple populations. *Physical Review D*, 91(2):023005, 2015.
- [83] N. Andres et al. Assessing the compact-binary merger candidates reported by the MBTA pipeline in the LIGO-Virgo O3 run: probability of astrophysical origin, classification, and associated uncertainties. *Classical and Quantum Gravity*, 39(5):055002, 2022.
- [84] B. Abbott et al. Search of S3 LIGO data for gravitational wave signals from spinning black hole and neutron star binary inspirals. *Physical Review D*, 78(4):042002, 2008.
- [85] C. Van Den Broeck et al. Template banks to search for compact binaries with spinning components in gravitational wave data. *Physical Review D*, 80:024009, 2009.
- [86] T. Dal Canton et al. Impact of precession on aligned-spin searches for neutron-star-black-hole binaries. *Physical Review D*, 91(6):062010, 2015.
- [87] B. P. Abbott et al. GW150914: First results from the search for binary black hole coalescence with advanced LIGO. *Physical Review D*, 93(12):122003, 2016.
- [88] D. A. Brown, I. Harry, A. Lundgren, and A. H. Nitz. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Physical Review D*, 86:084017, 2012.
- [89] I. W. Harry et al. A hierarchical search for gravitational waves from supermassive black hole binary mergers. *Classical and Quantum Gravity*, 25(18):184027, 2008.
- [90] S. Babak. Building a stochastic template bank for detecting massive black hole binaries. *Classical and Quantum Gravity*, 25:195011, 2008.
- [91] I. W. Harry et al. A stochastic template placement algorithm for gravitational wave data analysis. *Physical Review D*, 80:104014, 2009.

- [92] A. Nitz et al. Pycbc: Gravitational-wave data analysis toolkit, v2.3.3, 2024.
- [93] Pycbc template bank generation documentation (tmpltbank module). PyCBC official documentation. <https://pycbc.org/pycbc/latest/html/tmpltbank.html>.
- [94] T. Dal Canton et al. Implementing an aligned-spin search in pycbc. *Phys. Rev. D*, 90:082004, 2014.
- [95] C. Capano et al. Implementing a search for gravitational waves from non-precessing, spinning binary black holes. *Physical Review D*, 93:124007, 2016.
- [96] K. Kacanja et al. Efficient stochastic template bank using inner product inequalities. *The Astrophysical Journal*, 975:212, 2024.
- [97] S. Roy et al. Hybrid geometric-random template placement algorithm for gravitational wave searches from compact binary coalescences. *Physical Review D*, 95:104045, 2017.
- [98] S. Roy et al. Effectual template banks for upcoming compact binary searches in Advanced-LIGO and Virgo data. *Physical Review D*, 99:024048, 2019.
- [99] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. *Communication of the ACM*, 51:107–113, 2008.
- [100] Antoni Ramos-Buades et al. Next generation of accurate and efficient multipolar precessing-spin effective-one-body waveforms for binary black holes. *Physical Review D*, 108(12):124037, 2023.
- [101] E. Cuoco et al. Enhancing gravitational-wave science with machine learning. *Machine Learning: Science and Technology*, 2(1):011002, 2021.
- [102] E. Cuoco et al. Applications of machine learning in gravitational wave research with current interferometric detectors. *Living Reviews in Relativity*, 28:1, 2025.
- [103] A. Torres-Forné et al. Application of dictionary learning to denoise LIGO’s blip noise transients. *Physical Review D*, 102(2):023011, 2020.
- [104] M. Llorens-Monteagudo et al. Classification of gravitational-wave glitches via dictionary learning. *Classical Quantum Gravity*, 36:075005, 2019.
- [105] N. Mukund et al. Transient classification in LIGO data using difference boosting neural network. *Physical Review D*, 95:104059, 2017.
- [106] S. B. Coughlin et al. Classifying the unknown: discovering novel gravitational-wave detector glitches using similarity learning. *Physical Review D*, 99(8):082002, 2019.
- [107] R. Ormiston et al. Noise reduction in gravitational-wave data via deep learning. *Physical Review Research*, 2:033066, 2020.
- [108] M. Saleem et al. Demonstration of machine learning-assisted low-latency noise regression in gravitational wave detectors. *Classical and Quantum Gravity*, 41:195024, 2024.
- [109] T. Dooney et al. DeepExtractor: Time-domain reconstruction of signals and glitches in gravitational wave data with deep learning. *Physical Review D*, 112:044022, 2025.
- [110] M. Zevin et al. Gravity spy: Integrating advanced LIGO detector characterization, machine learning, and citizen science. *Classical and Quantum Gravity*, 34(6):064003, 2017.

- [111] Wu Y. et al. Advancing glitch classification in gravity spy: Multi-view fusion with attention-based machine learning for advanced LIGO's fourth observing run. *Classical and Quantum Gravity*, 42(16):165015, 2025.
- [112] M. Zevin et al. Gravity spy: lessons learned and a path forward. *The European Physical Journal Plus*, 139:100, 2024.
- [113] Calculation of a constant q spectral transform. *Journal of the Acoustical Society of America*, 89:425–, 01 1991.
- [114] F. Robinet. Omicron: An algorithm to detect and characterize transient noise in gravitational-wave detectors. Virgo Technical Document, 2016. VIRGO-TDS-14693.
- [115] F. Robinet et al. Omicron: A tool to characterize transient noise in gravitational-wave detectors. *SoftwareX*, 12:100620, 2020.
- [116] R. Essick et al. iDQ: Statistical inference of non-gaussian noise with auxiliary degrees of freedom in gravitational-wave detectors. *Mach. Learn.: Sci. Technol.*, 2:015004, 2021.
- [117] P. Godwin. Low-latency statistical data quality in the era of Multi-Messenger astronomy, 2020.
- [118] X. Ma, G. Vajente, et al. A deep learning technique to control the non-linear dynamics of a gravitational-wave interferometer. *Classical and Quantum Gravity*, 41:045003, 2024.
- [119] M. Coughlin et al. Limiting the effects of earthquakes on gravitational-wave interferometers. *Classical and Quantum Gravity*, 34(4):044004, 2017.
- [120] A. Biswas et al. New methods to assess and improve LIGO detector duty cycle. *Classical and Quantum Gravity*, 37(17):175008, 2020.
- [121] J. Buchli et al. Improving cosmological reach of a gravitational wave observatory using deep loop shaping. *Science*, 389(6764):1012–1015, 2025.
- [122] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [123] S. Markidis. The old and the new: Can physics-informed deep learning replace traditional linear solvers? *Frontiers in Big Data*, 4, 2021.
- [124] S. G. Rosofsky, H. Al Majed, and E. A. Huerta. Applications of physics informed neural operators. *Mach. Learn. Sci.*, 4:025022, 2023.
- [125] Sofia-Chrysovalantou Fragkouli et al. Deep residual error and bag-of-tricks learning for gravitational wave surrogate modeling. *Applied Soft Computing*, 147:110746, 2023.
- [126] P. Nousi et al. Autoencoder-driven spiral representation learning for gravitational wave surrogate modelling. *Neurocomputing*, 491:67–77, 2022.
- [127] N. Christensen et al. Parameter estimation with gravitational waves. *Reviews of Modern Physics*, 94(2):025001, 2022.
- [128] Gregory Ashton et al. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(1):27, 2019.
- [129] Ivan Kobyzev et al. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):3964–3979, 2021.

- [130] Gabriele Demasi et al. Enhancing gravitational-wave host localization with SKYFAST: rapid volume and inclination angle reconstruction. *Astronomy & Astrophysics*, 694:A245, 2025.
- [131] S. Babak et al. Searching for gravitational waves from binary coalescence. *Physical Review D*, 87(2):024033, 2013.
- [132] P. T. Baker et al. Multivariate classification with random forests for gravitational wave searches of black hole binary coalescence. *Physical Review D*, 91(6):062004, 2015.
- [133] S. J. Kapadia et al. Classifier for gravitational-wave inspiral signals in nonideal single-detector data. *Physical Review D*, 96(10):104015, 2017.
- [134] K. Kim et al. Ranking candidate signals with machine learning in low-latency searches for gravitational waves from compact binary mergers. *Physical Review D*, 101(8):083006, 2020.
- [135] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.
- [136] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, 1989.
- [137] Tanmaya Mishra et al. Optimization of model independent gravitational wave search using machine learning. *Physical Review D*, 104(2):023014, 2021.
- [138] V. Gayathri et al. Enhancing the sensitivity of transient gravitational wave searches with gaussian mixture models. *Physical Review D*, 102(10):104023, 2020.
- [139] M. J. Szczepańczyk et al. Search for gravitational-wave bursts in the third advanced ligo-virgo run with coherent waveburst enhanced by machine learning. *Physical Review D*, 107(6):062002, 2023.
- [140] Leigh Smith et al. Enhancing search pipelines for short gravitational wave transients with gaussian mixture modelling. *Physical Review D*, 110(8):083032, 2024.
- [141] Breiman L. et al. *Classification and Regression Trees*. Wadsworth, 1984. Originally published by Wadsworth, now frequently reprinted by Taylor & Francis / CRC Press.
- [142] J. Ross Quinlan et al. *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers, San Mateo, California, 1993.
- [143] J. L. Horowitz et al. Bootstrap methods in econometrics. *Annual Review of Economics*, 11:193–224, 2019.
- [144] Fabian Pedregosa et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [145] C. E. Metz et al. Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, 8(4):283–298, 1978. PMID: 112681.
- [146] Powers D. M. W. Evaluation: From precision, recall and f-measure to ROC, informedness, markedness & correlation. *Journal of Machine Learning Technologies*, 2(1):37–63, 2008.
- [147] Satyendra Singh Rawat, Amit Kumar Mishra, et al. Review of methods for handling class imbalance in classification problems. In *Handling Imbalanced Datasets in Machine Learning: Techniques and Applications*, volume 1038 of *Lecture Notes in Electrical Engineering*, pages 3–14. Springer Nature Singapore, 2024.

- [148] Chuan Guo et al. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70, pages 1321–1330, 2017.
- [149] A.H. Nitz et al. Detecting binary compact-object mergers with gravitational waves: Understanding and improving the sensitivity of the pycbc search. *The Astrophysical Journal*, 849(2):118, nov 2017.
- [150] B. P. Abbott et al. Binary black hole population properties inferred from the first and second observing runs of advanced LIGO and advanced virgo. *Astrophysical Journal Letters*, 882(2):L24, 2019.
- [151] B. P. Abbott et al. Supplement: The rate of binary black hole mergers inferred from advanced LIGO observations surrounding GW150914. *The Astrophysical Journal Letters*, 883(L1), 2016.
- [152] Murray Rosenblatt et al. Remarks on some nonparametric estimates of a density function. *Annals of Mathematical Statistics*, 27(3):832–837, 1956.
- [153] Gravitational wave open science center. <https://www.gwosc.org>.
- [154] Alexander H. Nitz et al. Catalog of gravitational waves from compact-binary mergers. *The Astrophysical Journal*, 946(2):59, 2023.
- [155] Marx E. et al. A machine-learning pipeline for real-time detection of gravitational waves from compact binary coalescences. *Physical Review D*, 111(4):042010, 2025.
- [156] Nousi P. et al. Deep residual networks for gravitational wave detection. *Physical Review D*, 108(2):024022, 2023.
- [157] V. Skliris et al. Toward real-time detection of unmodeled gravitational wave transients using convolutional neural networks. *Physical Review D*, 110(10):104034, 2024.
- [158] Alistair McLeod et al. Binary neutron star merger search pipeline powered by deep learning. *Physical Review D*, 111(2):024035, 2025.
- [159] Keiron O’Shea, Ryan Nash, et al. An introduction to convolutional neural networks. *arXiv preprint*, 2015.
- [160] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.
- [161] Marvin Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, Cambridge, MA, USA, 2017.
- [162] David E. Rumelhart et al. Learning representations by back-propagating errors. *Nature*, 323:533–536, 1986.
- [163] Suvrit Sra et al., editors. *Optimization for Machine Learning*. Neural Information Processing Series. The MIT Press, Cambridge, MA; London, England, 2011.
- [164] Kaiming He et al. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, June 2016.
- [165] B. P. Abbott et al. GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. *Physical Review D*, 93(12), 2016.

-
- [166] L. Pompili et al. Laying the foundation of the effective-one-body waveform models seobnr5: Improved accuracy and efficiency for spinning nonprecessing binary black holes. *Physical Review D*, 108(12), December 2023.
- [167] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Physical Review D*, 59(8), March 1999.
- [168] T. Dal Canton et al. Effect of sine-gaussian glitches on searches for binary coalescence. *Classical and Quantum Gravity*, 31(1):015016, 2014.
- [169] Brown D. A. et al. Detecting binary neutron star systems with spin in advanced gravitational-wave detectors. *Physical Review D*, 86(8), 2012.