Optimisation of Gravitational-Wave Data Analysis

Employing Machine Learning in Searches for Cosmic Strings and Binary Black Holes

Reinier Henricus Antonius Johannes (Quirijn) Meijer

ISBN: 978-90-393-7879-3
DOI: https://doi.org/10.33540/3019
Printed by Gildeprint – https://www.gildeprint.nl/
Cover design by Emmy Severijns
Based on Jean-Pierre Luminet's depiction of a black hole [1].

This work originates as part of the research program of the Foundation for Fundamental Research on Matter (FOM), and falls as of April 1, 2017 under the responsibility of the Foundation for Nederlandse Wetenschappelijk Onderzoek Instituten (NWO-I), which is part of the Dutch Research Council (NWO).

Optimisation of Gravitational-Wave Data Analysis

Optimalisatie van Data Analyse voor Zwaartekrachtgolven

(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht op gezag van de rector magnificus, prof. dr. ir. W. Hazeleger, ingevolge het besluit van het College voor Promoties in het openbaar te verdedigen op woensdag 27 augustus 2025 des middags te 2.15 uur

door

Reinier Henricus Antonius Johannes Meijer
geboren op 1 september 1992
te Oss

Promotor:
Prof. dr. C.F.F. Van Den Broeck

Co-promotor:
dr. S.E. Caudill

Assessment Committee:
Prof. dr. S.C.M. Bentvelsen
dr. A. Grelli
Prof. dr. S. Hild
dr. A. Samajdar
Prof. dr. S.J.G. Vandoren

Contents

Chapter 1

Introduction

Before starting the exposition to the theory of general relativity and gravitational waves, this introductory chapter serves as a buffer to build context for the physical ideas of gravity. As such, it frames the work in the following chapters, and was written in such a way that repetition between the chapters is minimised.

1.1 General Relativity

Gravity, from *gravitas* for weight, is the phenomenon of mutual attraction between masses. It is responsible for controlling the trajectories of astrophysical bodies at a macroscopic scale, such as planetary orbits, but its effects are evidently also found on Earth. Not only is it the gravity of Earth that ties humans to its surface, it is also the gravity of the Moon that is responsible for tidal cycles. As the Moon orbits Earth, its gravitational field will cause flood and ebb, or the respective rise and fall of water levels. Many other examples exist, and theories of gravity are therefore applicable in fields such as cosmology, astronomy, engineering and space travel. For an example of the latter, one may consider the gravity assist maneuver for spacecrafts, where the vehicle is assisted by the gravity of an astronomical object to alter its trajectory. This was famously executed using Jupiter and Saturn in the Voyager program [2]. In daily life a prime example of the applications of gravity comes in the form of the global positioning system (GPS) [3]. As the satellites that span the GPS network orbit Earth, their on-board atomic clocks will measure time at a rate that slightly differs from the rate that is measured by an identical clock on the surface of the Earth. This time dilation is due to the effects of gravity, with the satellites being further out in the gravitational field of Earth. The clocks therefore require time corrections, which are in fact computed using the theory of relativity, an advanced theory of gravity.

The motivation for relativity starts with a prior theory of gravity. Before general relativity, gravity was understood in terms of Newtonian mechanics, in what is called Newtonian gravity. Newton's universal law of gravitation states that gravity acts as a force [4]

$$F = G\,\frac{m_1 m_2}{r^2} \tag{1.1}$$

between two labeled objects with masses m_1 and m_2 that are separated by a distance r, and where G is the gravitational constant. At the most elementary level this law describes the mutual attraction between all pairs of particles, modelled on top of a model of time and space. Together with Newton's laws of motion, the law of gravitation can be used to describe the classical mechanics of particles that are subject to gravity. Although now superseded by relativity, this theory of gravity still holds as a simpler approximation in cases where gravity is weak and where speeds are slow compared to the speed of light.

Newtonian gravity stood until Maxwell introduced his highly succesful theory of electromagnetism. In particular, Maxwell's theory is incompatible with the structures of time and space assumed in Newtonian mechanics, as Maxwell's equations establish the speed of light as a universal constant, and the symmetries that electromagnetism obey are not those of Newtonian mechanics. However, particles that are both in motion and hold electromagnetic charges exist in the intersection of the two theories, meaning the two theories had to be reconciled. This objective is what lead Einstein to formulate special relativity in 1905 [5], giving a model of time and space together with a maximum speed in the form of the speed of light. This model can then serve as a basis for electromagnetic fields. In 1908, Minkowski gave a geometric formulation of this model that is now called a *spacetime* [6]. That a maximum speed is incompatible with Newtonian mechanics can be deduced directly from Newton's law of gravitation. If a particle were to be created and added, no matter at what distance the particle would be placed, the gravitational force between all other particles would be experienced immediately. This is what is called *instantaneous action*. In special relativity, where no information can travel faster than light, this can not be. In more detail, Newton's equations retain their form to all observers regardless of their inertial frame of reference, but Maxwell's equations do not. The symmetries of special relativity, the transformations between such frames, are bound by the speed of light. This shows both the importance of reference frames and the need to change the understanding of gravity to account for the permissible transformations between these frames.

Instead of trying to amend Newtonian gravity, Einstein published the theory of general relativity as an entirely new model for gravity in 1915 [7]. The choice for this replacement was motivated by two ideas. The first of these ideas is that of the *equivalence principle*, that states all bodies are affected by gravity, and all bodies fall in precisely the same way in a gravitational field. This is often explained through an observer that can not tell the difference between acceleration of his frame and gravity acting on his frame, establishing an equivalence between mass and gravitational mass. The motion of an object is therefore independent of its nature, and the paths of freely falling bodies define a preferred set of curves in spacetime. This suggests the axiom that the properties of a gravitational field are related to the structure of spacetime itself. The second idea follows from *Mach's principle*. The philosophy of Mach conjectures that in a universe without matter, concepts such as acceleration and rotation would hold no meaning. Therefore, everything is relative to matter, and matter should contribute to the definitions of these concepts. Einstein reasoned that a theory of gravity would therefore need to concern itself with the structure of spacetime, and involve the influence of matter. In general relativity this is precisely the case. Gravity manifests as the curvature of spacetime, and not as a force acting within spacetime. In converse, the curvature of spacetime is determined by the distribution of matter.

General relativity was first succesfully tested by Eddington in 1919 [8]. However, the tests of general relativity done before the first direct detection of gravitational waves in 2015, treated in the next subsection, were limited. At current, the accepted model of cosmology is entirely relativistic, taking the form of the Friedmann-Lemaître-Robertson-Walker model [9]. Despite this it is known that the theory of general relativity is not yet complete, as the problem of gravity at quantum scales remains open.

1.2 Gravitational Waves

Gravitational waves are displacements in gravitational fields, generated by the motion of mass. Since they do not have strong interactions with matter in the way that electromagnetic waves do, they are unique messengers for astronomical information. For example light can be deflected and absorbed, but a gravitational wave does not face these obstructions, with gravitational lensing [10] being the exception. Gravitational waves will therefore contain information that has been mostly unmodified from the moment of emission, enabling the precise observation of exotic objects that are extremely far from Earth.

Although general relativity is the theory that envelops gravitational waves, the first predictions of their existence predate relativity. In 1893, Heaviside proposed gravitational waves as a counterpart to electromagnetic waves [11]. The leading theory of gravity at that time was Newtonian gravity however, which does not allow for wave propagation because of instantaneous action. Any information would teleport more so than propagate in a wave. The first prediction with relativistic speed constraints was made by Poincaré in 1905 [12] following the requirement formulated by Lorentz, that said all forces would have to obey the symmetries of special relativity. This put gravity and electromagnetism on equal stature, and both should have waves. As the acceleration of electrical charge produces electromagnetic waves, so too must accelerated mass in a gravitational field produce gravitational waves. In 1916, Einstein demonstrated that gravitational waves do in fact result from general relativity [13]. Einstein himself however doubted his own predictions, and in 1936 submitted a paper claiming gravitational waves could not exist. Einstein had convinced himself that gravitational waves must be remnants of frame transformations.

The first evidence for gravitational waves, albeit indirect, was found after 1974 [14, 15]. During observation of the orbital progression of the Hulse-Taylor binary pulsar it was found that its decay follows the predictions of general relativity, from which it can be calculated that a certain amount of energy should be lost to gravitational radiation. This amount is in agreement with the expected orbital decay and won Hulse and Taylor the 1993 Nobel Prize in Physics. The idea of gravitational waves carrying and transmitting energy hinges on the *sticky bead argument* proposed by Feynman in 1957. This argument is a thought experiment designed to show the physical effects of a passing gravitational wave. If a rod with two freely sliding beads on it is considered, and a gravitational wave were to pass over this rod, the length of the rod would practically remain fixed due to other natural forces, but the distance between the beads should oscillate. As the beads move over the rod, the friction must generate heat, and thus energy is being transferred. Detailed versions of this argument were published by Bondi and Pirani in the same year [16]. In a separate sequence of events, Weiss recognised that this argument could be used to measure the transmitted energy using lasers. This would lay the groundwork for the first direct observation of a gravitational wave.

Around the same time Weiss obtained this insight, Thorne worked on the theoretical development of gravitational waves, and Barish was gaining experience in other experimental high-energy physics projects, before the three took on key roles in the LIGO collaboration. In 2015, the centennial year of general relativity, the detectors of the LIGO-Virgo collaboration made their first direct observation of a gravitational wave. This wave was produced by the coalescence of two black holes, and the analysis of this event was presented to the world in a 2016 publication [17]. The following year Weiss, Thorne and Barish would be awarded the 2017 Nobel Prize in Physics for their contributions. In response, Thorne has been quoted as saying:

"It is unfortunate that, due to the statutes of the Nobel Foundation, the prize has to go to no more than three people, when our marvelous discovery is the work of more than a thousand."

This quote not only showcases the immense scale of the experiment, but also further underlines the difficulty of the design. The contributors persisted when even the father of the theory was doubtful, thinking that even if possible, the measurement of gravitational wave amplitudes would forever be outside of human reach.

1.3 Black Holes

Black holes are stars that famously speak to the imagination, often making appearances in popular media. Even before being captured in picture accurate depictions existed [1], and artist renditions existed in media. A notable example is the fictional black hole Gargantua in the 2014 film Interstellar. This depiction was supported by Thorne, who both published the calculations leading to this visualisation and wrote a popular science exposition for the physics that lay at their foundation [18].

In 2019 the Event Horizon Telescope generated the data for the first picture of a black hole called Messier 87* [19]. This was a landmark achievement, showing the presence of a dark hole through its shadow. This picture was sharpened in 2023 by re-rendering it on the same data, but using an improved algorithm [20]. The James Webb Telescope is currently observing black holes in the optical and infrared bands, which will possibly lead to new visualisations.

Apparent from the previous subsection, black holes are examples of moving masses that generate gravitational waves. Contrary to what is sometimes believed, black holes are not giant, insatiable objects that seek to eat mass. They will however grow by absorbing mass that gets caught in their extreme gravitational attraction. This leads to an increase in mass of the black hole, and can lead to black holes of extreme mass called supermassive black holes. The consensus is that such black holes exist at the centers of most galaxies, with one called Sagittarius A* in fact existing in the center of the Milky Way.

In theory black holes, like gravitational waves, also arise from general relativity, where they are predicted as regions of spacetime where gravity is so strong that not even light can escape. At the very center of these regions all mass of the black hole will be contained, with density and gravity that will both tend to infinity. The border of the region is demarked by the event horizon, described by Finkelstein as a perfect unidirectional membrane. Information on events past the horizon and inwards to the black hole region cannot reach outside observers. The black hole therefore traps everything, and is black in the sense that in relativistic theory it is devoid of light, and therefore of both colour and information.

The appearance of infinities and black holes signal the breakdown of general relativity to some. Others, however, argue that at such scales the effects should be described by quantum theory, believing the theory to only be incomplete. As mentioned before in this chapter, gravity and quantum theory have not been unified to date, making black holes a central point to the study of quantum gravity.

Again as was mentioned previously, black holes admit no visible features according to general relativity. Most observation would therefore have to happen indirectly, through for instance the gravitational effects on passing light in the form of lensing, or the tracing of the paths of bodies that pass nearby a black hole. Alternatively, black holes may form rings of matter called accretion disks. These rings are formed from accreted matter that circles a black hole as it falls inwards, accelerating in its trajectory and heating through friction. The heat in turn leads to electromagnetic radiation in for instance the X-ray band. Such radiation is what enabled the first indirect detection of a black hole called Cygnus X-1 in 1971 [21].

Stars, and therefore black holes, will often come with companion stars in binary systems. Such systems are formed when the stars are bound in the orbit of one another through gravity. This requirement of gravitational interaction is what sets them apart from optical doubles that are only required to visually appear close together. Once a body is exposed to the tremendous gravitational pull of a black hole, it is likely the body will end up being absorbed. Other black holes are no exception. Thus, binary black holes are systems of two black holes on course to absorb each other in a merger following their inspiral. If this happens the result is a single black hole that is more massive than the two individual black holes, with part of their masses lost to the emission of gravitational waves. Binary black holes are one of the strongest known sources of gravitational waves, followed by binary neutron stars and neutron star-black hole binaries.

1.4 Cosmic Strings

Cosmic strings, in brief, are hypothetical objects that manifest as massively dense strings on cosmic scales. Compared to black holes these objects are a lot more elusive. Where black holes were predicted by general relativity, cosmic strings are predicted in theories of particle physics. Their gravitational effects are however predicted by general relativity. They were first contemplated by Kibble in 1976 [22], but to this day no observations have been made.

In a way, cosmic strings can be considered to be cracks in the universe, arising as a side-effect of a phase transition. This formulation presents a way of picturing their formation mechanism. Consider a body of water that is freezing, undergoing a phase transition to ice. It may occur that the water will start freezing at separate points, and as the water continues to freeze fully, these frozen sections will inevitably meet. At the interface of these sections imperfections will remain as defects of the medium. Cosmic strings would then be the imperfections that form and grow as the universe evolves and expands.

During the evolution of the universe the cosmic strings should evolve into a network of cosmic strings, consisting in particular of closed loops. All such closed loops are necessarily doomed. As dynamic objects, they would oscillate and radiate away energy in the form of gravitational waves, eventually leading to their evaporation. Furthermore, strong concentrated emissions of gravitational waves can happen at several points on the string. Examples of such points are the intersections of interacting cosmic strings, or the formation of cusps and kinks.

At one point in time it was thought that the gravity of cosmic strings might have been responsible for the formation of galactic superclusters. However, it was later calculated that their contribution to structure formation in the universe would be at most 10% [23]. This led to a declining interest in their study, that was reversed during what is called the second superstring revolution in string theory. At this time different string theories were unified, and string-theoretical research exploded. Although cosmic strings are not exclusive to string theory, string theory is notoriously plagued by a lack of observable signatures, and cosmic strings might prove to be one. Currently no such results have been found. Instead, null-results are used to bound their parameters and give indication of where to look, and where not to look.

1.5 Machine learning

Machine learning is a subset of artificial intelligence, specifically containing machinal algorithms that learn and are able to perform tasks that are beyond the tasks fed to the algorithms in the training regime. This is the driving force of machine learning, as many current applications show incredible generalising power. It is also the difference between machine learning and traditional statistics. Machine learning focuses on the previous objective, the finding of patterns that generalise. Statistics on the other hand focuses on the computation of statistics from data samples for the purpose of inference. In a sense, machine learning is the outwards directed complement to the inwardly turned statistics. Although the conceptualisation and development of modern machine learning started in the last century, it has only been in recent years that computing power and data storage capabilities have reached a point where the learning processes can be accomodated. In the meanwhile, there have been two artificial intelligence winters (in the timeframes of 1973-1980 and 1987-2000, although there is no consensus on the precise years), times during which the interest in artificial intelligence cooled down, with resources leaving the field. As of right now, machine learning is seeing wide application in fields with difficult problems, such as natural language processing, computer vision and audio recognition.

Two types of machine learning algorithms are central to this thesis.

The first are neural networks, algorithms that are constructed from single computational units called neurons, mirroring the biological neurons in the brain. In modelling a network of neurons many degrees of freedom remain, such as the weight or importance assigned to information that is passed through the network. The idea at the core of the modern neural network is that few such decisions are made in the design of the network, and these values are tuned based on data that is fed to the network in order to learn a task. The dataset is then an incomplete sample of the full information on the task, and the neural network will attempt to learn the task by adjusting its weights. If done correctly, the hope is that the model will generalise to the full unseen information.

The second type is that of evolutionary algorithms, also called genetic algorithms. Evolutionary algorithms are applicable when an optimisation problem can not be formulated in such a way that an efficient algorithmic plan exists to find its optima. In that case, an evolutionary algorithm can be designed, that will find a solution as a complicated and stochastic extension of bisection methods, inspired by natural evolution. Solutions to the optimisation problem are collected in a generation of solutions, and every following generation of solutions is determined by the best solutions from the previous generation. The hope would then be that the quality of the solutions will eventually stabilise near an optimum, and therefore ideally give at least an approximation to a global optimum. A famous example where evolutionary algorithms were applied is in the evolved antenna [24].

1.6 Contributions

The work in this thesis is based on three papers,

- Gravitational-Wave Searches for Cosmic String Cusps in Einstein Telescope Data using Deep Learning (published in Physical Review D) [25];
- Robustness of Deep Learning Models to Precession in Gravitational-Wave Searches for Intermediate-Mass Black Hole Binaries [26] (submitted to Physical Review D) and;
- Fast Waveform Generation for Gravitational Waves using Evolutionary Algorithms (published in Physical Review D) [27],

with each corresponding to one of the respective Chapters 4, 5 and 6.

For cosmic string searches, a neural network was implemented as a binary classifier. This model learned to distinguish between specific hypothetical gravitational-wave signals originating from cosmic strings and structural detector noise called glitches that can greatly resemble such signals. This work was done in the context of a future-generation gravitational-wave detector, and therefore also serves as a science case. The reason for this is that current-generation detectors have made no cosmic string detections, and the proposed detector will have a greatly increased resolution in the frequency band where the signals are expected to occur. Because this model accurately learned the classification task, and in fact greatly outperformed the current state-of-the-art in cosmic string searches, work was done to interpret the inner workings of the model in an attempt to learn what it is that the model has learned.

The same neural network was retrained for a specific mass class of black hole binaries as they would be observed by the current generation of detectors. Searches for this class of black hole binaries are not only hindered by their likeness to glitches, as cosmic string searches are, but are also complicated by additional effects that are unique to binary systems, such as the dynamics of the orbital planes. In addition, the specification of such systems require a lot more parameters than do cosmic strings. It was shown that the performance of the model is highly effective for this problem as well, again outperforming the state-of-the-art, and the model was studied within the context of the high-dimensional parameter space for binary black hole waveforms. This allowed for the deeper understanding of the model's performance in gravitational-wave searches.

The third paper studies the potential increase in effiency of waveform models using evolutionary algorithms, specifically applied to the waveform generation for binary black holes. In the current implementation, waveforms are generated by taking a frequency spectrum and computing the power of the waveform at every individual frequency in this spectrum. This is a costly procedure, and it may not always be necessary to compute these powers for every frequency. Instead, one can attempt to compute the exact value at a subset of the frequency spectrum, and interpolate the values in the complement. The problem of where in the frequency band the exact values are required and where interpolation will suffice is an optimisation problem, studied through the application of evolutionary algorithms.

Chapter 2

Physics

In this chapter two major but as of yet irreconcilable theories in modern physics are introduced in a succint fashion: general relativity, the theory of the very large and very massive, from which black holes and gravitational waves stem, and quantum field theory, the theory of the smallest quanta, that introduces cosmic strings. The theoretical basis of both is required to fully motivate the discussion in this thesis.

2.1 General Relativity

Following the high-level overview in the introduction, this section delves into general relativity in detail. This exposition differs from typical treatments that lead into gravitational waves in that the use of index notation is kept to a minimum. The underlying philosophy is that general relativity can be fitted onto the geometric theory, much like a dialect is a linguistical variety.

The majority of the material in this section is inspired by [10, 28–34], and these books are referenced throughout.

2.1.1 Spacetimes and Gravity

For the entirety of this section, M will be used to denote a smooth manifold with or without boundary [35]. Recall that a Riemannian metric is defined as follows.

Definition 2.1.1 (Riemannian metric)**.** A Riemannian metric g on M is a smooth, symmetric and covariant 2-tensor field that is positive definite at each point.

The pair (M, g) is called a *Riemannian manifold*. Riemannian manifolds can be classified by their signature, a basis independent invariant of the metric. For any matrix representing g, the signature is an ordered triple (p, q, r) of the number of positive, negative and zero eigenvalues respectively, all counted with multiplicity. General relativity, however, requires a more general structure that is defined similarly to the Riemannian manifold, but with a relaxation on the requirement of positive definiteness.

Definition 2.1.2 (Pseudo-Riemannian metric)**.** A pseudo-Riemannian metric g on M is a smooth, symmetric and covariant 2-tensor field that is non-degenerate at each point, and has the same signature everywhere.

Instead of positive definiteness, the metric of a *pseudo-Riemannian manifold* (M, g) is guaranteed non-degeneracy and a constant signature, which in turn guarantees that the fundamental theorem of Riemannian geometry extends to pseudo-Riemannian manifolds. This implies that like Riemannian manifolds, a pseudo-Riemannian manifold admits the unique torsion-free connection called the Levi-Civita connection [29]. It is for this reason the pseudo-Riemannian manifold is considered a generalisation for which the same notation can be used. Note however that whereas every smooth manifold has a Riemannian metric, this result does not necessarily hold true for strictly pseudo-Riemannian metrics [29].

An important special case is that of the pseudo-Riemannian manifolds with metric signature $(n-1, 1, 0)$, called *Lorentzian manifolds*. This specific signature allows for a local time orientation when a Lorentzian manifold is taken to be a model for time and space. The inner products induced on the tangent bundle by the metric allow for the separation of three regions within the tangent spaces, classified through the norms of these inner products. The first region is a double cone defined by the *null vectors* of the norm, which span the two *light cones* connected at their vertex in the origin. The geometry of this subspace can be deduced from diagonalisation of the matrix representing the metric, possible due to the assumption of symmetry. The interior vectors of the light cones are those vectors for which the norm is negative, called the *timelike vectors*. By defining an equivalence relation on the vectors in the tangent space where two vectors x, y are related if and only $g(x, y) < 0$, two equivalence classes coinciding with the interiors of the light cones arise. By designating one to be *future-directed* and the other to be *past-directed*, the time orientation, and therefore a causal structure, has been constructed. The naming conventions stem from the fact that timelike vectors are only separated by time, with the intuition being that one can move from one timelike vector to another in the future direction as time progresses. In general relativity this translates to nothing traveling faster than the speed of light, as photons are assumed to travel on null geodesics following the tangent null vectors [31]. The remainder of the vectors in the exterior of the light cones are *spacelike vectors*, separated by space, and therefore unreachable from points on or within the light cones unless the speed limit is violated. Such a model, including the time orientation, will be referred to as a *spacetime*.

An element of a spacetime is called an *event*, and an event is per definition characterised by the $n-1$ spatial dimensions and a time. Related to the photons carrying light, an example of an event might be the presence of a particle at a certain spatial location at a given time. In extension, as time evolves, so too might the spatial location of the particle. As such, the particle traces a continuous curve parametrised by time which is called the *worldline* of the particle. These are examples of how spacetimes can be used to model a variety of processes depending on time and space, both static and dynamic.

From the pragmatic point of view of a geometer, relativity then is the analysis of the Lorentzian manifolds used to model spacetimes. For a physicist however, relativity concerns itself with the back and forth translations between spacetime models and physical observations.

The first example one usually encounters is that of the Minkowski spacetime $\mathbb{R}^{3,1} = (\mathbb{R}^4, \eta)$ that characterises special relativity, with its metric η defined by the diagonal matrix:

$$\begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \tag{2.1}$$

The Minkowski spacetime is a flat Lorentzian 4-manifold wherein, due to the signature, the first dimension is reserved for time, and the remaining three are reserved for physical space. Its flatness can be deduced from $\mathbb{R}^{3,1}$ being locally isometric to $\mathbb{R}^4$. As such, its Riemann tensor vanishes everywhere. For a spacetime, this lack of curvature is equivalent to the absence of gravity. This is a postulate that will be asserted in more detail later in this section. Minkowski spacetime has the practical advantage that every tangent space is isomorphic to the Minkowski space itself [31], so that light cones can be visualised as parts of the spacetime itself. This makes many structures intuitively accessible, and it is because of these properties that the special case of Minkowski spacetime is studied in special relativity.

The Minkowski metric specifies the *spacetime interval* of special relativity, giving a measure of the separation in time and space. For the difference vector $\Delta x \in \mathbb{R}^{3,1}$ between two fixed events, this quantity is given by:

$$\sum_{\mu,\nu=0}^{3} \eta_{\mu\nu} \mathrm{d}x^\mu \otimes \mathrm{d}x^\nu (\Delta x, \Delta x) = -\Delta x_0^2 + \Delta x_1^2 + \Delta x_2^2 + \Delta x_3^2. \tag{2.2}$$

The one-forms in this equation are obtained through the musical isomorphisms, also called the *raising* and *lowering* of indices in Ricci calculus. Labeling the coordinates as (t, x, y, z), the literature often presents this as a differential line element in the following expression:

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + \mathrm{d}x^2 + \mathrm{d}y^2 + \mathrm{d}z^2 \tag{2.3}$$

and moreover requires this quantity to be invariant under permissible coordinate transformations between different *frames.* These transformations are given by the Lorentz group $\mathrm{O}(3,1)$ consisting of those linear transformations that leave the origin fixed, namely the three-dimensional rotations of space, and boosts that effect time through simultaneity in different frames. More generally, the Lorentz group is the isotropy subgroup of the Poincaré group $\mathbb{R}^{3,1} \rtimes \mathrm{O}(3,1)$ which is the full isometry group of $\mathbb{R}^{3,1}$ that also includes affine transformations through the semi-direct product.

Moving beyond the special case of Minkowski spacetime, curvature of time and space is introduced in general relativity. The physical theory is that this curvature stands in direct relationship with the distribution of matter. This relationship is caught in a *field equation*, a partial differential equation that determines the dynamics of a tensor field, usually representing some physical measurement. Einstein postulated the following field equation for the class of spacetimes:

$$E + \Lambda g = \kappa T. \tag{2.4}$$

In the left hand side of this equation, $E = \mathrm{Ric} - \frac{1}{2}Rg$ is a $(0,2)$-tensor called the *Einstein tensor*, defined using the Ricci curvature tensor Ric and the scalar curvature R. On the right hand side, the $(0,2)$-tensor T is the *stress-energy-momentum tensor* that quantifies the local stress, energy and momentum as attributes of matter. The remaining variables $\Lambda, \kappa \in \mathbb{R}$ are the *cosmological constant* and *Einstein gravitational constant* respectively. Although the cosmological constant is usually set to 0, including it as a non-zero parameter will later enable intrinsically curved models of the universe, with famous examples being de Sitter and anti-de Sitter spaces [33]. For the Einstein gravitational constant, it follows from the full expression:

$$\kappa = \frac{8\pi G}{c^4} \tag{2.5}$$

that the fixing of the natural units at $G = c = 1$ leads to the value $\kappa = 8\pi$. Including the gravitational constant ensures that in the appropriate limit, general relativity reduces to Newtonian gravity. The extremely small value of the constant, driven by the denominator, is also informative of the incredible density required to induce even a small degree of curvature. Following this discussion the Einstein field equation can be recast into the following simple form:

$$E = 8\pi T. \tag{2.6}$$

The key element that ties gravity to the curvature described by this equation is the postulate that the worldlines of free falling objects are geodesics. Hence, gravity becomes an inherent part of the structure of spacetime, summarised by John Wheeler in the following quote [36]:

"Matter tells spacetime how to curve, and spacetime tells matter how to move."

A solution to the Einstein field equation is a metric tensor, or per extension, a spacetime. In the special case where T is set identically to zero, the equation is referred to as a *source-free* or *vacuum* equation, and the spacetime solution is called a *vacuum solution*. One should note that the absence of a source does not imply the absence of curvature in the solution, as is the case with for example the famous Schwarzschild solution (treated in subsection 2.1.5). The vanishing of the Ricci tensor in the vacuum equation does not imply the vanishing of the Riemann tensor. As it is only the tensor trace of the Riemann tensor, the Ricci tensor does not hold complete information on the curvature of the manifold.

Finding analytical solutions to the Einstein field equation is non-trivial. Because the expansion of the Ricci tensor depends non-linearly on the Christoffel symbols [29], and the Christoffel symbols involve the metric, the Einstein field equation is equivalent to a coupled system of non-linear second order partial differential equations in the metric. Solving partial differential equations, or even proving the existence of solutions, is a famously daunting task [37,38], as is evidenced by the Navier-Stokes problem's inclusion in the Millennium Problems [39].

In brief, the prior statement about the contents of relativity can be rephrased for geometers and physicists alike as follows: general relativity studies the solutions of the Einstein field equations and their properties.

2.1.2 Gravitational Waves

Specifically because of the difficulty of solving the Einstein equation, many methods are studied to reduce the complexity of this problem. Such a reduction may take the shape of certain assumptions on the solution's symmetry, or even a move away from finding exact solutions to finding non-exact solutions by way of approximation. It is the latter method that is studied in this section. In accordance with the maxim of perturbation theory, it will be assumed the solution is a perturbation of the flat Minkowski metric. This assumption means the curvature of the spacetime solution would be limited, or equivalently, that gravity is weak in this model. In practice this leads to a linear field theory that has proven capable of handling fluctuations of the metric.

As an initial condition, assume that the metric tensor of the spacetime can be decomposed into the Minkowski metric plus a deviation:

$$g = \eta + h, \tag{2.7}$$

for h small. Small, in this context, is understood to mean the determinant of the matrix defining h is bound by a fixed value δ. The above decomposition can then be substituted for g in the Einstein field equation, after which all non-linear terms are discarded. This results in the linearised Einstein equation [10,31], signified by the order in the superscript:

$$E^{(1)}(h) = 8\pi T. \tag{2.8}$$

Relativity is a *gauge theory*, meaning there are several degrees of freedom retained within the linearised equation, also referred to as *gauge freedom*. In the case of relativity this freedom represents a type of diffeomorphism invariance, capturing the notion that there exist vector fields on the manifold so that the perturbation

h satisfies the linearised equation if and only if $h + \mathcal{L}_V g$ does [31,33], where $\mathcal{L}_V g$ denotes the Lie derivative of g with respect to $V \in \mathfrak{X}(M)$. It follows from their definition that Killing fields form particular examples. In relativity the addition of such a term is called fixing a *gauge*, and this freedom generalises to the full Einstein field equation. In the current context of the linearised Einstein equation it is possible to choose a frame and gauges so that the equation reduces to a wave equation [10,31]:

$$\Box h = -16\pi T \tag{2.9}$$

where $\Box$ denotes the d'Alembertian operator. Discussion of the specific choices of gauges and frame is suspended until the next subsection. This equation admits two solutions of the form:

$$h_p = H_p \exp\left(i \sum_{\mu=0}^{3} k_\mu x^\mu\right) \tag{2.10}$$

for $p \in \{1,2\}$. The H_p in this expression that distinguish the two solutions are constant $(0,2)$-tensor fields, and k is the constant wave vector the direction and norm of which determine the direction of propagation and wavelength respectively. The perturbation of the flat Minkowski metric, or *gravitational wave*, h can now be expressed as a linear combination of these two solutions over $\mathbb{R}$:

$$h = h_+ h_1 + h_\times h_2. \tag{2.11}$$

It is customary to use the plus and cross symbols for the labeling of these solutions, owing this to the geometric interpretation of the associated tensors. This will be discussed further in the next subsection. The tensors are called the polarisation modes, whereas their coefficients are called the polarisation coefficients.

As a note on their propagation, gravitational waves travel at the speed of light. This follows from the fact that the wave vector should be a null vector so that the h_p satisfy the wave equation, where the speed of light enters into the d'Alembertian, together with the previous discussion of null vectors defining light cones. Per transitivity, gravitational waves travel at the speed of light.

Summarising the past discussion, gravitational waves are waves that can be considered to travel across spacetime, and are therefore given shape as a disturbance of the Minkowski background metric. The fluctuations are determined by linearising the Einstein field equation, which leads to a two-dimensional solution space, spanned by basis elements that have distinct polarisations. Because general relativity directly couples spacetime curvature to gravity, these waves are interpreted as waves of gravity. The remainder of this section will be spent on a very brief review of four types of gravitational waves. Since the astrophysical objects that generate these gravitational waves have not been introduced, the paragraphs are written in such a way that a deep understanding is not required.

The first type is that of *continuous gravitational waves* [40]. A spinning neutron star can be idealised as being a perfect sphere spinning around an axis at a constant speed. If an irregularity is added to the surface of this sphere, the mass distribution will see a time-dependent change as the star revolves around its axis, which causes the continuous emission of a gravitational wave. In this ideal situation the frequency of the passing waves will be constant.

In a sense a *burst gravitational wave* [41] is the opposite of a continuous gravitational wave. Rather than being periodic, these waves are characterised by a single short-lived burst. An example of an event generating a burst gravitational wave is a supernova (from *super stella nova*), the explosion of a star. As the star explodes outwards in all directions, this redistribution of matter causes a burst gravitational wave.

The preceding two types of gravitational waves could be demonstrated through single objects. The type of *compact binary coalescence gravitational waves* [42] require a binary system of two objects. Compact in this context refers to the objects having high density due to a high mass relative to their ratio, and has no connection to the topological notion of compactness. Black holes are the archetypical examples of compact objects, and a binary black hole system can be used to exemplify this type. The lifecycle of the binary will involve the decaying orbit of the two black holes, until they finally coalesce. In orbit this binary system will generate gravitational waves that increase in frequency with the orbital frequency, until the violent merger into a single black hole that generates the peak intensity of the gravitational wave. These systems are studied in detail in subsection 2.1.5.

The fourth and final type is that of *stochastic gravitational waves* [43]. Rather than being generated by single events that can be isolated, a stochastic gravitational wave is a combination of many gravitational waves that have melded together over time. Stochastic gravitational waves are sometimes collectively referred to as *the gravitational-wave background* as a counterpart to the cosmic microwave background. The defining property of this background is that it can not be analysed by imposing the same patterns as are present in the other three types, and is therefore studied through stochastic properties under the assumption of randomness instead. It is conjectured that beyond astrophysical sources, a large contribution is made by cosmological sources from the early universe.

2.1.3 Interactions of Gravitational Waves with Mass

Having theoretically motivated the existence of gravitational waves as well as giving a cursory explanation of their generating mechanisms, a logical next venture would be the treatment of their experimental observation. In order to do so, it is required to build an understanding of how gravitational waves interact with measurable quantities within a spacetime.

Starting from the linearised Einstein field equation in Eq. 2.8, it is required to fix both the gauges and a frame to arrive at the wave equation in Eq. 2.9. This will fix eight out of the ten degrees of freedom in the linearised equation, forced by the matrix symmetry, leaving the polarisations. The first course of action is rewriting the linearised equation in terms of the trace-reversed metric perturbation h':

$$h' = h - \frac{1}{2}\mathrm{tr}_\eta(h)\eta, \tag{2.12}$$

which is allowed since the deviations from h are lost in the higher order terms, with the additional assumption that in the weak-field limit the Minowski metric can be used for the trace-reversal [10]. Note the similarity of this procedure to obtaining the Einstein tensor as the trace-reversed Ricci tensor. The change in parity of the trace eliminates a number of terms in the expansion of the Einstein tensor. The other remaining terms, save for $\Box h'$, are then eliminated by fixing the harmonic gauge [31] defined as the vector field solution V to:

$$\Box V = -\partial^\nu h'_{\mu\nu}. \tag{2.13}$$

to arrive at Eq. 2.9 with six degrees of freedom. The name of this gauge originates from its relation to Laplace's equation and its solutions as it forces $\partial^\nu h'_{\mu\nu} = 0$, and is also called the de Donder gauge or Lorenz gauge. Interestingly, the latter is often erroneously written as Lorentz, stemming from the similarity of the names of Lorenz and Lorentz.

More degrees of freedom can be fixed, as gauge freedom remains under the condition that the corresponding vector field is chosen from the kernel of the d'Alembertian, so that it does not interfere with the harmonic gauge. This freedom is taken by fixing the transverse-traceless gauge [10,31] in which h' becomes transverse through the requirement $h'_{0\nu} = 0$, and traceless by requiring $\mathrm{tr}_\eta(h') = 0$. It follows from the traceless condition that within this gauge h must coincide with h'.

With both gauges fixed, far removed from the generating system the source-free linearised Einstein equation will have the two independent solutions given in Eq. 2.10. Choosing a frame so that the transverse wave aligns with the x_3-axis or z-axis, and taking the real part that corresponds to physical solutions so that only the cosine term remains, the general solution in Eq. 2.11 can be expressed as:

$$h = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \cos[\omega(t - z)]. \tag{2.14}$$

The structure of the matrix can be read from the conditions set prior. Per transversality and symmetry the waves must be purely spatial and limited to only

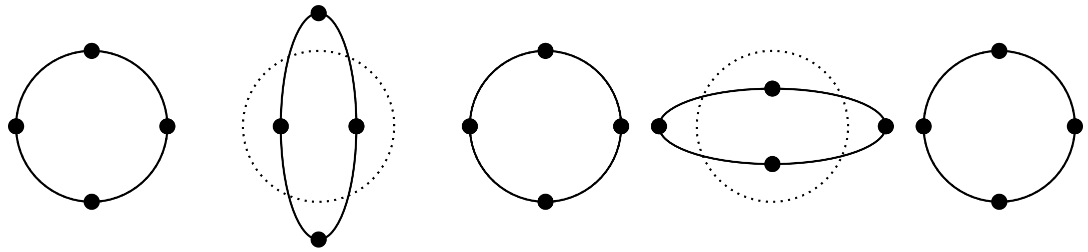

Figure 2.1: The interaction of a plus-polarised gravitational wave passing transversely through the xy-plane, stretching and compressing the particle ring within.

affecting the coordinates x_1 and x_2, meaning $h_{0\nu} = h_{\mu 0} = 0$ and $h_{3\nu} = h_{\mu 3} = 0$. Furthermore, it follows from the traceless condition that $h_{11} = -h_{22}$, and finally from the symmetry requirement for Lorentzian metrics that $h_{12} = h_{21}$. Since g was assumed Lorentzian in Eq. 2.7 and η is symmetric, h must retain this symmetry in order to not violate the initial assumption. This leaves two degrees of freedom for the polarisations, labeled h_+ for h_{11} and $h_\times$ for h_{12} respectively.

Substituting the above expression Eq. 2.14 into the decomposition Eq. 2.7, the segment $\mathrm{d}s^2$ can be read off in (t, x, y, z)-coordinates as:

$$\mathrm{d}s^2 = -\mathrm{d}t^2 + (1 + h_+ \cos[\phi])\mathrm{d}x^2 + (1 - h_+ \cos[\phi])\mathrm{d}y^2 + 2h_\times \cos[\phi]\mathrm{d}x\mathrm{d}y + \mathrm{d}z^2 \quad (2.15)$$

with $\phi := \phi(t) = \omega(t - z)$. The behaviour of the two polarisations can be understood from this form by setting the opposing polarisation identical to zero. Doing so for $h_\times$ to consider only the plus-polarised solution, the physical distance between two points separated by a distance $\mathrm{d}x$ along the x-axis is given by:

$$\mathrm{d}s = \sqrt{(1 + h_+ \cos[\phi])\mathrm{d}x^2} = \sqrt{1 + h_+ \cos[\phi]}\,\mathrm{d}x \approx (1 + \frac{1}{2} h_+ \cos[\phi])\mathrm{d}x \quad (2.16)$$

using the Babylonian hypothenuse approximation $\sqrt{a^2 + b^2} \approx a + b^2/(2a)$, and similarly for y by:

$$\mathrm{d}s = \sqrt{1 - h_+ \cos[\phi]\mathrm{d}y^2} \approx (1 - \frac{1}{2} h_+ \cos[\phi])\mathrm{d}y. \quad (2.17)$$

Considering a ring of particles in the xy-plane, it can be read from the inclusion of the cosines and cosine parities that a passing plus-polarised wave must repeatedly stretch and compress the particle ring. This is illustrated in Fig. 2.1. For the cross-polarisation similar behaviour can be deduced by setting $h_+ = 0$ and noting that the resulting expressions are similar to the ones found above under a 45° rotation of the xy coordinate system.

2.1.4 Gravitational-Wave Observatories

A number of the properties of gravitational waves discussed thus far can be utilised to design experimental observatories, and more broadly, networks of such observatories. It was previously discussed that passing gravitational waves affect the relative separation between particles, and this forms the foundation of the interferometric gravitational-wave detectors.

Interferometers measure the patterns created by the recombination of a separated light source, which reveal information on how the light has traveled between its source and the recombination point. Recombination in this context happens according to the superposition principle. Light waves behave linearly, meaning the waves are additive and homogeneous. If two such light waves are superimposed, the individual amplitudes can be summed, and the result is called the interference pattern. In this subsection the Michelson interferometer will be introduced.

The Michelson interferometer was originally designed for the Michelson-Morley experiment meant to detect the velocity of Earth in the luminiferous aether, which was believed to exist and be the medium for light waves at the time of the experiment [44]. Despite this experiment disproving their theory, the configuration has proven to be highly influential, inspiring both the postulate that the speed of light is a universal constant as well as forming the basis for both current and future proposed gravitational-wave detectors. At a minimum, a *Michelson interferometer* consists of a light source, a beam-splitting mirror, two reflective mirrors and a photodiode that converts absorbed photons to electrical currents. The light source emits a laser beam directed at the beam splitter, from which the beam is separated into two separate beams that are reflected back off the mirrors on their paths before being recombined in the central photodiode. A measurement of the output current results in data called the *detector strain* or relative deformation.

An assumption that is incorporated into this configuration is that upon recombination at the photodiode, the light waves are out of phase by half a wavelength, so that destructive interference takes place. This is equivalent to a strain measurement of zero under ideal circumstances. If however a gravitational wave should pass, the interference patterns may change, as the paths traveled by the beams are contracted and elongated, in turn leading to constructive interference and non-zero measurements. The differential changes can be derived in the same way as the effects on the particle ring were, assuming that the arms that encase the laser beams are aligned with the axes of the xy-plane and the gravitational wave propagates normal to this plane. Since photons follow null geodesics, one can set $\mathrm{d}s^2 = 0$ in Eq. 2.15 and derive the following effect on the arm lengths for a purely plus-polarised gravitational wave:

$$\mathrm{d}x = (1 + h_+ \cos[\phi])^{-1/2}\mathrm{d}t, \quad \mathrm{d}y = (1 - h_+ \cos[\phi])^{-1/2}\mathrm{d}t. \tag{2.18}$$

These expressions show the periodic oscillations in time. Repeating this with the same assumptions for a purely cross-polarised wave demonstrates the possible blindness of a detector, and therefore the need for a network of detectors with different orientations. The vertices in the network graph are then connected through edges with weights depending on the speed of light, as the time of arrival in a different detector can be computed from the distance and propagation speed. Besides increasing visibility, the building of networks is useful for the sky localisation of sources, as well as coincidence through cross-correlation.

Currently operational interferometers use designs that are more complex. In particular, it was suggested in the discussion of Eq. 2.5 that gravitational waves detected at a large distance from their source must have extremely small amplitudes. This translates to measuring arm length differentials of approximately 10^{-5} the width of a proton, currently resulting in gravitational-wave strains with order of magnitude 10^{-21}. Measurement of such small values requires very long arms. The advantage in sensitivity is unbounded seeing as when the length of the arms increases, an absolute differential will increase proportionally, meaning there is a linear relationship between arm length and detector sensitivity. For this reason modern designs implement *Fabry-Peról cavities* that artificially extend arm lengths [45]. Such a cavity is realised by inserting an additional mirror on the laser trajectory immediately after the beam splitter, where the newly added mirror forces the light to remain within the cavity through resonance. Through this construction the laser beam can be recycled within the cavity, forcing a longer path.

A result of this high level of sensitivity is the registration of unwanted signals or *noise* [46]. There exist several classes of noise, with examples being seismic noise from the Earth, thermal noise in the mirrors, gas noise in the vacuum laser enclosures, and electronic noise from the measurements. These examples are non-exhaustive, and a specific additional class of unwanted signal is that of *glitches*. Glitches are transient artefacts that often mimic astrophysical signals [46], making their mitigation a crucial part of detector data analysis. Despite being typically unmodelled, their recurrence allows for the separation of glitches into different types, usually based on their duration and frequency ranges.

The current and second generation of gravitational-wave detectors includes two locations of LIGO (Laser Interferometer Gravitational-Wave Observatory) [49], with one being located in Hanford (Washington, US) and another in Livingston (Louisiana, US). Both of the dual detectors are operated by the LSC (LIGO Scientific Collaboration) and share their design, with arm lengths of 4 kilometers. In Europe, the Virgo detector in Santo Stefano a Macerate (Cascina, Italy) [50] is operated by the EGO (European Gravitational Observatory) consortium. This detector is shown in Fig. 2.2a. The design of the Virgo detector resembles that of the two LIGO detectors, with major differences including arm lengths of 3 kilometers each, along with an alternative mirror-suspension structure. A second European detector called the GEO600 detector [51] is maintained by the GEO collaboration

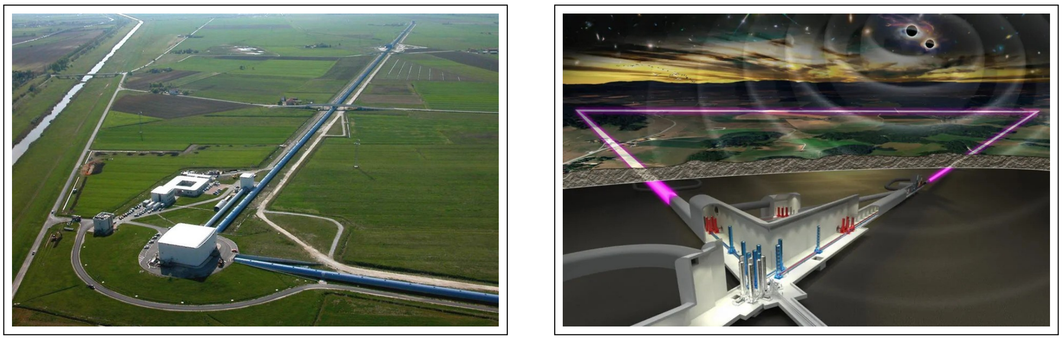

(a) The Virgo observatory in Italy. (b) An artistic rendition of the future Einstein Telescope.

Figure 2.2: Visualisations of both current second-generation and future third-generation gravitational-wave observatories, taken from [47] and [48] respectively.

in Sarstedt (Lower Saxony, Germany), a collaboration of institutes in Germany and the United Kingdom. The name of this detector stems from its arm lengths at 600 meters. Finally, KAGRA (Kamioka Gravitational Wave Detector) [52] was built in Asia, in Hida (Gifu prefecture, Japan). The design of this detector differs dramatically from that of the aforementioned detectors, as it is placed underground to combat the seismic noise from Japan's notable earthquake activity, and uses a cryogenic system to cool its mirrors and reduce thermal noise. Under the umbrella of the LVK (LIGO-Virgo-KAGRA) collaboration, these observatories carry out joint data analyses with co-authored results papers. The detectors in this network alternate between *observational phases* where data is taken, and *scientific phases* during which technical improvements are implemented.

Once a set of proposed improvements is large enough, or a new observatory in line with the standards of such upgrades is built, one can say the product enters a new generation. In the current second generation, the LIGO and Virgo detectors were upgraded from the first generation to aLIGO (Advanced LIGO) and aVirgo (Advanced Virgo). Plans for the third generation have been proposed, with for example Cosmic Explorer [53]. Cosmic Explorer is set to follow a similar design to the current generation of detectors, with arm lengths of 40 kilometers. Another proposed third-generation detector is the underground Einstein Telescope [54,55]. The Einstein Telescope will consist of three detectors with arms of 10 kilometers in length, configured in an equilateral triangle. The space-based LISA (Laser Interferometer Space Antenna) observatory [56, 57] is to share this configuration, as it is designed as a system of three spacecrafts forming an equilaterial triangle that will trail the Earth in heliocentric orbit. Similar to the Einstein Telescope, this formation creates a triangular observatory with arm lengths of 2.5 million kilometers. An artist rendition of the triangular design of the Einstein Telescope is shown in Fig. 2.2b.

Observatories are not limited to laser interferometry. An example would be

NANOGrav (North American Nanohertz Observatory of Gravitational Waves) [58] as part of the International Pulsar Timing Array, that rely on pulsar timing arrays as their basis for observation. A pulsar, from *pulsating radio source*, is a magnetised and rotating neutron star that emits electromagnetic radiation from its poles. By the extreme density of neutron stars, their rotation has a regular frequency that strictly defines the interval between pulses. This regularity can be used when an array of galactic pulsars is timed to form a pulsar timing array, and deviations in the arrival times on Earth can be used to measure low-frequency gravitational waves. As interferometers and pulsar timing arrays both share the electromagnetic spectrum, this concept mirrors that of interferometry.

As a final note, the combination of methodologies and fields allows for *multi-messenger astronomy*, the coordinated effort of astronomy through several parallel channels or *messengers*. Messengers beside gravitational waves include electromagnetic waves, cosmic rays, and neutrinos. Combining the information delivered by different messengers can greatly improve the understanding of astrophysical processes [59].

2.1.5 Compact Binary Coalescence

Previously, compact binary systems were mentioned as examples of sources of gravitational waves. In order to prepare for the analysis of such gravitational waves as they would be present in detector data, this subsection covers the theory of compact binaries, with a strong emphasis on black holes.

As both black holes and neutron stars are stellar remnants, their study starts with stars. Stars are astronomical objects of study, defined as luminous spheroids of plasma held together by self-gravity. The mass of a star is measured in solar masses (with one *solar mass* $M_\odot$ corresponding to the mass of the Sun, circa $1.99 \cdot 10^{30}$ kg), and stars are called *compact* if their mass is sufficiently high relative to their radius, meaning the star passes a density treshold. This treshold is typically set in comparison to non-compact stars. Stars evolve throughout their lifetime, in what is called the process of stellar evolution, dictated in large part by the mass of the original star. In particular, when the mass exceeds the Tolman-Oppenheimer-Volkoff limit of approximately 2.5 solar masses, the inwards acting gravitational force will eventually become so overwhelming that no other natural force can stand up to it. At this stage the mass is compressed until such a density is reached that the star essentially becomes a point-mass, the strongest possible example following the compactness definition, as the density now approaches infinity. This qualitative infinity is related to a quantitative infinity in the geometry of general relativity: *black holes* coincide with points in spacetime where the curvature tends to infinity, also called *spacetime singularities*. Alternatively, singularities can be defined through the geodesic incompleteness of the manifold [31, 32].

The first gravitational field of a black hole was found by analytically solving the

Einstein field equation. Karl Schwarzschild derived what is now known as the *Schwarzschild solution*, also called the Schwarzschild metric [60]:

$$\mathrm{d}s^2 = -\left(1 - \frac{r_s}{r}\right)\mathrm{d}t^2 + \left(1 - \frac{r_s}{r}\right)^{-1}\mathrm{d}r^2 + r^2\mathrm{d}\Omega^2, \tag{2.19}$$

given in the spherical Schwarzschild coordinate system [10]. Due to the assumptions necessary to find this solution, the Schwarzschild black hole is the simplest example of a black hole. It is assumed to be spherically symmetric, static, and non-rotating. The metric admits one physical parameter $r_s = 2Gm$ that depends on the gravitational constant G and the object mass m. This quantity is called the *Schwarzschild radius*, and corresponds to the distance of the singularity to the spherical event horizon that surrounds it. The *event horizon* is the boundary of communication. Beyond the event horizon and towards the singularity, the gravitational potential well becomes so strong that not even light can cross back over the horizon. This is the reason for the black hole nomenclature, as black is the absence of light. The Schwarzschild radius had previously been computed for dark stars in Newtonian gravity, by equating the speed of light to the escape velocity that would be needed to break free from a spherically symmetric massive body, in casu the dark star. This formulation results in a different physical interpretation. In Newtonian gravity, a particle would be able to break free from the dark star, but eventually fall back onto the surface. In general relativity, the particle can never move outwards past the event horizon of the black hole.

One should note that the Schwarzschild solution evidences the existence of *coordinate singularities*, which differ from spacetime singularities, like the singularity at $r = 0$, in that they can be removed by a coordinate transformation. Specifically by transforming to Eddington-Finkelstein coordinates, the coordinate singularity at the event horizon, where $r = r_s$, can be removed [10].

By the no-hair theorem [36], a static black hole is in the limit uniquely characterised by three parameters, being its mass, its spin in the form of the angular momentum, and its electromagnetic charge. This theorem bears no relation to the hedgehog theorem, also called the hairy ball theorem [61]. Leaving the mass a free parameter, this leads to four types of black holes, with the non-spinning and chargeless *Schwarzschild black hole* being the trivial example. In addition, charged but non-rotating black holes are called *Reissner-Nordström* black holes, *Kerr-Newman black holes* both rotate and carry charge, and the *Kerr black holes* rotate but lack charge [62]. As will become apparent in this subsection, Kerr black holes that correspond to Roy Kerr's solution [63]

$$\begin{aligned}\mathrm{d}s^2 = &-\left[1 - \frac{2rm}{\rho^2}\right]\mathrm{d}t^2 - \frac{2rma\sin^2(\theta)}{\rho^2}(\mathrm{d}t\mathrm{d}\phi + \mathrm{d}\phi\mathrm{d}t) + \frac{\rho^2}{\Delta}\mathrm{d}r^2 + \\ &\rho^2\mathrm{d}\theta^2 + \frac{\sin^2(\theta)}{\rho^2}\left[(r^2+a^2)^2 - a^2\Delta\sin^2(\theta)\right]\mathrm{d}\phi^2\end{aligned} \tag{2.20}$$

in Boyer-Lindquist coordinates [10] are of particular importance to compact binary coalescences. This coordinate system is a generalisation of the Schwarzschild coordinates, as in the zero-spin limit the Schwarzschild coordinates are recovered.

Addressing the previously free mass parameter, black holes can be classified into mass ranges, which coincidentally determine their size as measured through the Schwarzschild radius. The *stellar-mass black holes* occupy the range of $10^0 - 10^1 M_\odot$, *supermassive black holes* cover the range of $10^5 - 10^9 M_\odot$, and *intermediate-mass black holes* span the intermediate mass range. Other named classifications exist, with ultramassive black holes being those black holes with mass upwards of the supermassive black holes, and micro black holes (also called quantum mechanical black holes) [64, 65] that are sub-solar-mass theoretical curiosities. The interplay of quantum mechanics and gravity will be touched upon again in the next chapter.

An alternative terminating point for stellar evolution, other than a black hole, is a *neutron star*. If the mass of the star remnant does not exceed the Tolman-Oppenheimer-Volkoff limit [66, 67], but is above the Chandrasekhar limit of approximately 1.4 solar masses [68], the star will eventually undergo the process of core-collapse in a supernova, with a neutron star as its remnant. Stars with masses below this limit will evolve into white dwarves, which are not treated here. The neutron star that remains will retain most of its angular momentum, causing it to spin, and over time will undergo spin-down in isolation. Neutron stars, like black holes, attain extreme densities. As such, neutron stars too have a Schwarzschild radius, albeit smaller than the diameter of the star. This is a notable departure, as this fact enables light to escape the star and be astronomically observed in the electromagnetic spectrum. Without this, pulsar timing arrays could not be.

A key observation in the context of gravitational waves is that not all black holes and neutron stars form through stellar evolution. It is for instance hypothesised that *primordial black holes* have been created through the dense packing of subatomic matter in the early universe [69], but of more importance to this thesis, compact objects can coalesce and merge in a compact binary coalescence [34]. The remainder of this subsection will be dedicated to this process and its dynamics.

Compact objects may come in a binary system, wherein the objects are captured in mutual orbit. This process of *compact binary coalescence* knows three phases [34]. The first is the *inspiral*, during which time the objects move in orbit. This movement produces gravitational waves, and in order to compensate for the energy lost through this radiation, the orbital radius reduces, and the orbital frequency must increase. A back-reaction is therefore in effect during the inspiral, where energy is increasingly radiated away at the cost of the radius, up until the innermost stable circular orbit is reached. At this point the orbit can no longer be sustained, and the objects plunge into a *merger*. The result of the merger phase is a singular perturbed Kerr black hole that converges to a stable configuration

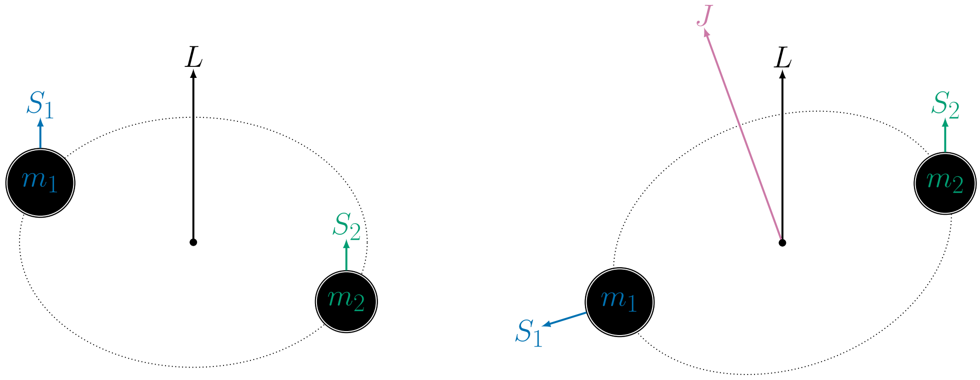

(a) A binary system with aligned spins. (b) A binary system with misaligned spins.

Figure 2.3: Depictions of two compact binary systems. The spin vectors of the objects in the system on the left align, meaning the total angular momentum aligns with the orbital angular momentum. In the system on the right, the spin vectors are misaligned, and therefore a non-zero angle between the orbital angular momentum and total angular momentum is introduced.

during the *ringdown* phase.

In the scenario where one object is a black hole and the other is a neutron star, additional mechanics are activated. As the mass of the neutron star is not concentrated to a point, and tidal forces are present, the mass of the neutron star will undergo tidal deformation in a region called the Roche lobe of the black hole [70]. In this region the mass of the neutron star is bound to the black hole, and mass can in fact be transferred to the black hole in the process of Roche lobe overflow. Note that this is one of the mechanisms through which the black hole might form an *accretion disk*, a disk of accreted diffuse material that moves around the central black hole. This can be extended to the case of a binary neutron star, where the deformation of both neutron stars have to be considered.

As the example of mass transfers illustrate, the current presentation of compact binary coalescence depends on a degree of idealisation. In practice the orbits are not required to be circular, and may in fact be eccentric [34]. A different possibility is that the orbital plane defined by the two objects along a fixed axis is not stable, causing precession of the orbital plane in the same way a spinning gyroscope might precess [34]. This phenomenon is the focus of the next paragraph.

Recall that the *orbital angular momentum* L of the system is defined as the cross-product of the position vector and the linear momentum. As such, the orbital angular momentum will be perpendicular to the plane spanned by these two vectors. It is a typical choice to align the z-axis of the source frame with L. Define the *total angular momentum* J as the sum $J = L + S$ of the orbital angular momentum and the total spin vector S. If J and L are aligned, and conversely no spin angular momentum is present, the source system is called *aligned* or *anti-*

aligned, depending on the sign of the spin vectors. Such a configuration is shown in Fig. 2.3a. If this is not the case however, the source system is called *precessing*. The result of adding the dynamical spin vector to L is an angle between L and J at which J traces a cone through time, called the *precession cone*. This development can be seen from Fig. 2.3b. Assuming that energy is mainly radiated along the direction of L due to the waves being transverse [34], one can see that the effect of orbital precession on gravitational waves is manifested in the form of modulations in amplitude and phase [71]. The analysis of gravitational waves can be simplified by ignoring the possibility of precession, but the assumption of alignment on the spin vectors is very strong, and will not hold true in general. At the time of writing, precession is increasingly included in studies.

Although this subsection has concerned itself with the review of compact binary sources of gravitational waves, it harbors the unary source of supernovae. These stellar explosions were previously mentioned in the context of burst gravitational waves, and this type of source will see a return when cosmic strings are treated in the next section.

2.1.6 Detector Data and Waveform Models

With the dynamics of compact binary coalescences established, the objective of this subsection is the parametrisation and modelling of the signals as they would appear in the data that is registered by a detector. Such a model is called a *waveform*.

In practice the raw strain from the interferometer is measured as an electric current. In theory, the data is obtained through the application of a constant $(2,0)$-tensor D, called the *detector tensor*, at discrete sample times. The data can therefore be written as a time series:

$$(Dh[t] : t \in \mathbb{N}) , \tag{2.21}$$

where $h(t)$ is the measured perturbation of the Minkowski metric at time t. The detector tensor is typically omitted and the strain time series is denoted by $s(t)$, with the Fourier transform $\tilde{s}(f)$ defining a corresponding frequency series.

The parameter space $\mathcal{P}$ of compact binary coalescences can be identified with $\mathbb{R}^{15}$. Of these 15 dimensions, two are reserved for the primary mass m_1 and secondary mass m_2, with the convention that the primary mass is taken to be the largest of the two. These parameters are supplemented by six spin parameters that together form two three-dimensional spin vectors S_1 and S_2 for the primary and secondary masses respectively. There are then four coordinates for the spacetime location. For time, the time of coalescence t_c in GPS time is taken for reference. The spatial coordinates are given by the luminosity distance d_l (giving the Euclidean distance taking into account the flux for a given luminosity) and sky location defined by the right ascension ϕ and declination θ. These angles are parallel to the longitude and latitude in geographical coordinates. Relating the detector frame to the source frame, the polarisation angle ψ is defined as the angle between the detector plane and the vector starting at the origin of this plane that ends at the source, and the inclination ι is the angle between the normal vector to the detector plane and the orbital angular momentum. The inclination is therefore a measure of the expected visibility in the detector. Finally, the last parameter is the orbital phase of coalescence ϕ_c.

The parameters in $\mathcal{P}$ can be split into two types. The first type is that of the *intrinsic parameters* (m_1, m_2, S_1, S_2) that describe the physics intrinsic to a system. The second type is that of the *extrinsic parameters* $(t_c, d_l, \phi, \theta, \psi, \iota, \phi_c)$ that relate the observing detector to the source. For these parameters, a note is made at the inclination. Following the direction of strongest gravitational-wave emission, a system with $\iota \in \{0, \pi\}$ is said to be *face-on*, as it follows from the discussion in the previous subsections that this is the ideal for gravitational-wave detection. The *edge-on* systems, for which $\iota = \pi/2$, per complement describe a worst-case detection scenario.

Given the extrinsic parameters, one can move between the source frame and the detector frame using coordinate transforms. Given a signal h with polarisations h_+ and $h_\times$ in the source frame, application of the *detector patterns* defined as:

$$F_+(\theta, \phi, \psi) = -\frac{1}{2}[1 + \cos^2(\theta)] \cos(2\phi) \cos(2\psi) - \cos(\theta) \sin(2\phi) \sin(2\psi) \tag{2.22}$$

and

$$F_\times(\theta, \phi, \psi) = \frac{1}{2}[1 + \cos^2(\theta)] \cos(2\phi) \sin(2\psi) - \cos(\theta) \sin(2\phi) \cos(2\psi) \tag{2.23}$$

to the individual polarisations by:

$$h(t) = F_+(\theta, \phi, \psi)\, h_+(t) + F_\times(\theta, \phi, \psi)\, h_\times(t) \tag{2.24}$$

will yield the signal in the detector frame. Effectively, this transforms the gravitational wave in the source frame to its registration in the detector frame.

For general systems, the parameters in $\mathcal{P}$ as defined here are not exhaustive. For binary systems that include neutron stars, parameters that account for the tidal deformability of the mass of the stars should be included, per the discussion in the previous subsection. As an additional example, a parameter that quantifies the deviation from circular orbit should be included when eccentricity is taken into account. These parameters are neglected here.

There are several ways in which the parameters in $\mathcal{P}$ can be recombined, given that the recombination is invertible. One example is a mapping of the masses called the *chirp mass*, which to leading-order largely determines the orbital evolution during the inspiral, and is typically easier to measure from data than the individual masses are [72]. By itself, the chirp mass

$$M_c = (m_1 m_2)^{3/5}(m_1 + m_2)^{-1/5} \tag{2.25}$$

is reductive, as this function can not be injective per symmetry. Including the *mass ratio* $q = m_2/m_1$ in $(m_1, m_2) \mapsto (M_c, q)$ creates a bijective map, as can be confirmed algebraically. Given the chirp mass by itself, the frequency evolution of an inspiral can be approximated by solving a given differential equation [34], and the name is chosen due to the increase in pitch of this solution. The masses can also be used to define the *dimensionless spin vectors*:

$$\chi_i = S_i/m_i^2. \tag{2.26}$$

The dimensionless spins can be used to measure the complementary effective spin and effective precessing spin parameters. The *effective spin*, inspired by the previous treatment of aligned and anti-aligned systems, describes the mass-weighted spins aligned with the orbital angular momentum:

$$\chi_{\text{eff}} = \frac{m_1\chi_1 + m_2\chi_2}{m_1 + m_2} \cdot L, \tag{2.27}$$

where $\cdot$ is the dot product. For a precessing system this quantity will not hold full information, as only the aligned spin components are considered. Therefore, with $\sigma_i = \cos^{-1}(\hat{L} \cdot \hat{S}_i)$ denoting the angles between the normalised orbital angular momentum and spin angular momenta, the *effective precessing spin*:

$$\chi_{\text{p}} = \max\left[|\chi_1|\sin(\sigma_1), \left(\frac{3q + 4q^2}{4 + 3q}\right)|\chi_2|\sin(\sigma_2)\right] \tag{2.28}$$

quantifies the in-plane spin components orthogonal to the orbital angular momentum as a measure of precession [73]. The asymmetry in the arguments follows from the assumption that the primary mass will always be the largest of the two.

Before moving to waveform models and their genealogy, the concept of higher-order modes is treated. Doing this before starting on waveforms will allow for a continuous treatment of the different models.

Like Lebesgue integrable functions can be decomposed into sums of circular functions through their Fourier transforms, so too can complex-valued spherical functions on $S^2 \subset \mathbb{R}^3$ be decomposed. In this construction the role of the exponential functions are taken over by the spherical harmonics, which form a complete and orthonormal basis for this function space [74, 75]. This basis forms a subspace of the $\mathcal{C}^2(\mathbb{R}^3, \mathbb{C})$-solutions of Laplace's equation,

$$\nabla^2 f(x, y, z) = 0, \tag{2.29}$$

which are collectively called the *harmonic functions*. The name of the *spherical* *harmonics* is therefore explained by the restriction of these functions to spheres. Laplace's equation is solved by transforming to spherical coordinates and using separation of variables to r and (θ, ϕ), leading to the spherical harmonics $Y_{\ell,m} : S^2 \to \mathbb{C}$ of degree ℓ and of order m with corresponding Legendre polynomials $P_{\ell,m} : [-1, 1] \to \mathbb{R}$. In this basis the general solution can be expressed as:

$$f(r, \theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} \left[f_{1,\ell,m} r^{\ell} + f_{2,\ell,m} r^{-\ell-1}\right] Y_{\ell,m}(\theta, \phi), \tag{2.30}$$

where the $f_{i,l,m} \in \mathbb{C}$ are coefficients. In the case of gravitational-wave emission, one can introduce time-dependent coefficients $h_{\ell,m}$ (also called *modes*) at a fixed radius that depend on the extrinsic parameters Θ of the system so that h can be decomposed as [34, 76]:

$$h(t, \Theta, \theta, \phi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} h_{\ell,m}(t, \Theta) Y_{\ell,m}^{(-2)}(\theta, \phi), \tag{2.31}$$

where $Y_{\ell,m}^{(-2)}(\theta, \phi)$ are the spin-weighted spherical harmonics of weight -2, constructed through the application of a lowering ladder operator to the spherical harmonics [77]. Similar to taking partial sums of Taylor series, the expansion of h can be approximated to leading orders by truncating the above sum. Assuming the system is face-on, this is typically done by taking the quadrupole mode corresponding to $(\ell, m) = (2, 2)$, as for such systems $Y_{2,2}^{(-2)}$ will attain a global maximum. If higher-order terms are included, this is colloquially referred to as the inclusion of *higher-order modes*. Depending on the extrinsic parameters, these additional terms can add a large contribution. For edge-on systems for example, $Y_{2,2}^{(-2)}$ will attain a local minimum, whereas other spherical harmonics will attain their global maxima.

Waveform models follow a tree structure. The most accurate solutions possible are those obtained from directly solving the Einstein field equation numerically, as is done in numerical relativity [78, 79]. This is incredibly computationally expensive, and has therefore lead to the collection of such solutions in catalogs for reference (see for instance [80–82]). Since commercial machines are not powerful enough to run these computations at the time of writing, practical applications that require waveforms for different systems rely on *waveform approximants*. These models are designed to approximate the numerical solutions to a reasonable accuracy, but at far greater speeds. At the same time, this implies the numerical solutions still play a vital role, as the catalogs are required for the calibration of the approximant models.

At the first node down from the numerical solutions, one finds the family of *surrogates* [83–85]. These are approximants that interpolate between the waveforms for known solutions in the catalogs. In addition, the family tree contains a set of descendent sibling nodes in the approximant families that directly model the two-body system evolution based on sets of assumptions, each with their own use-cases and regions of validity. The `Taylor` family [86, 87] uses post-Newtonian approximations that model solutions to the Einstein field equation in the weak-field limit by accounting for the deviations from Newtonian gravity. The approximants are therefore applicable to systems with a high mass ratio and with relatively low compactness. In practice these approximants are typically used to model the inspiral of non-spinning and non-precessing binary black holes. Following a different formalism, the effective-one-body (`EOB`) family reduces the two-body dynamics to the case of a corresponding single body [88,89]. As a result of this approach, these approximants are more effective for low mass ratio systems. Finally, there is the phenomenological family [90–95] that is based on the assumption that during the frequency evolution, the behaviour can be modelled through analytical closed-form expressions. This family is designed with computational speed in mind, and is therefore often employed when a large number of waveforms needs to be generated.

The latter `Phenom` family is used in this thesis, with all models originating in the derivation of the frequency-domain `IMRPhenomD` model [95], which models the quadrupole mode in the absence of precession. Factoring this mode into amplitude and phase functions, the assumption that

$$\tilde{h}_{2,2}(f, \Theta) \approx A(f, \Theta) e^{i\Phi(f,\Theta)}, \tag{2.32}$$

is made, where different functions for the amplitude and phase are substituted depending on the frequency range of the compact binary coalescence. The frequency band is separated into the inspiral, a connecting intermediate region, and the combined merger and ringdown, for a total of three disjunct ranges. These functions are then joined on the range boundaries using the step function $I_F(f)$ for $F \subseteq \mathbb{R}$, that equals 1 if and only if $f \in F$, and 0 otherwise.

The starting point Φ' of the inspiral phase is taken from post-Newtonian theory and shared with the `TaylorF2`-approximant [87], giving the phase of a non-spinning or aligned-spin point-particle. In compact binary systems however, the objects will have interacting spins, which are incorporated into the inspiral phase to obtain the following expression:

$$\Phi_{\mathrm{I}}(f, \Theta) = \Phi'(f, \Theta) + \frac{(m_1 + m_2)^2}{m_1 m_2} \left(\alpha_0 + \alpha_1 f + \frac{3}{4}\alpha_2 f^{4/3} + \frac{3}{5}\alpha_3 f^{5/3} + \frac{1}{2}\alpha_4 f^2 \right). \tag{2.33}$$

Here, the α_i-weighted terms are higher-order post-Newtonian terms which are calibrated using the numerical solution catalogs and `EOB`-approximants. In the connecting region the following observationally obtained closed form is used [95]:

$$\Phi_{\mathrm{C}}(f, \Theta) = \frac{(m_1 + m_2)^2}{m_1 m_2} \left(\beta_0 + \beta_1 f + \beta_2 \ln(f) - \frac{\beta}{3} f^{-3} \right). \tag{2.34}$$

The coefficients β_i in this expression are determined by the boundary conditions between the inspiral and merger-ringdown regions. In the merger-ringdown region, the expression:

$$\Phi_{\mathrm{MR}} = \frac{(m_1+m_2)^2}{m_1 m_2}\left[\gamma_0 + \gamma_1 f - \gamma_2 f^{-1} + \frac{4}{3}\gamma_3 f^{3/4} + \gamma_4 \arctan\left(\frac{f-\gamma_5 f_r}{f_d}\right)\right] \tag{2.35}$$

is used, again with calibrated coefficients γ_i. Other quantities of note in this formula are the ringdown frequency f_r of the remnant black hole and the ringdown damping frequency f_d. Both of these frequencies originate in the modelling of the ringdown as a series of quasi-normal modes [34].

The full expression for the phase is then:

$$\Phi_{\mathrm{IMR}}(f,\Theta) = I_{(-\infty,f_0]}(f)\cdot\Phi_{\mathrm{I}}(f,\Theta) + I_{(f_0,f_1)}(f)\cdot\Phi_{\mathrm{C}}(f,\Theta) + I_{[f_1,\infty)}(f)\cdot\Phi_{\mathrm{MR}}(f,\Theta), \tag{2.36}$$

where the f_i denote the transition frequencies between the three different regions.

Simlarly for the amplitude, the starting point A' is obtained from post-Newtonian theory, and the amplitude for the inspiral is given by:

$$A_{\mathrm{I}}(f,\Theta) = A'(f,\Theta) + \left(\frac{m_1+m_2}{\sqrt{m_1 m_2}}\right)\sum_{i=1}^{3}\delta_i f^{(5+2i)/6} \tag{2.37}$$

where the δ_i weigh terms in the same way the α_i-coefficients did for the phase. Again, similarly for the intermediate connecting region, the amplitude is given by:

$$A_{\mathrm{C}}(f,\Theta) = \frac{(m_1+m_2)^2}{m_1 m_2}\left(\epsilon_0 + \epsilon_1 f + \epsilon_2 f^2 + \epsilon_3 f^3 + \epsilon_4 f^4\right) \tag{2.38}$$

with the ϵ_i-coefficients determined by boundary conditions, and finally, the amplitude during the merger-ringdown phase is given by:

$$A_{\mathrm{MR}}(f,\Theta) = \frac{m_1 m_2}{(m_1+m_2)^2}\cdot\frac{\zeta_1\zeta_3 f_d}{(f-f_r)^2+(\zeta_3 f_d)^2}\cdot\exp\left[\frac{-\zeta_2(f-f_r)}{\zeta_3 f_d}\right] \tag{2.39}$$

with the ζ_i-coefficients calibrated against the numerical solutions. This leads to the amplitude:

$$A_{\mathrm{IMR}}(f,\Theta) = I_{(-\infty,f_0]}(f)\cdot A_{\mathrm{I}}(f,\Theta) + I_{(f_0,f_1)}(f)\cdot A_{\mathrm{C}}(f,\Theta) + I_{[f_1,\infty)}(f)\cdot A_{\mathrm{MR}}(f,\Theta). \tag{2.40}$$

Substituting Φ_{IMR} and A_{IMR} into Eq. 2.32 completes the `IMRPhenomD` model.

This approximant was updated to `IMRPhenomXAS` [92], short for generation `X` of the aligned-spin approximant in the series of `Phenom` waveform models. In turn, this model was extended to `IMRPhenomXHM` [91] that includes the higher-order modes for $(\ell, m) = (2,1), (3,3), (3,2), (4,4)$ by a procedure similar to the one above, and this approximant itself was upgraded to include precession in `IMRPhenomXPHM` [90]. Precession is included through what is called the twisting-up of the waveform, explained below.

A key observation is that gravitational waves from precessing binary systems are simplified when observed in a frame that adapts to the precessional motion [90]. Instead of aligning the z-axis of the frame with the orbital angular momentum, the frame can be dynamically transformed to track the total angular momentum. Doing so maps the precessing system onto an equivalent aligned-spin system, as the spin component parallel to the orbital angular momentum is what largely drives the precessional motion of the orbital plane [90]. Thus, a precessing system can be modelled by generating the corresponding aligned-spin waveform before applying the inverse rotation map, given by a set of time-dependent angles. These angles can be described by the Euler angles [90, 96], and effecting the frame transformations is called the *twisting-up* of the waveform.

The approximants that have been described are all frequency-domain approximants. An example of a time-domain approximant is `IMRPhenomTP` [93], which is a time-domain parallel to `IMRPhenomD` that was twisted up. The non-precessing counterpart `IMRPhenomT` [94] exists, but was introduced afterwards.

As a final note, after this subsection, the terms waveform and waveform approximant will be used interchangeably if no confusion can arise. In order to remain decoupled, the treatment of data analysis methods involving the waveforms that were introduced here is postponed to a later chapter.

2.2 Classical and Quantum Field Theory

Whereas general relativity lends itself to the study of massive objects, and quantum mechanics is famously used for the very small atomic and subatomic scales, classical mechanics works as an approximation in the interface of the theories. In order to transition from general relativity and to explain the reasoning behind the construction of quantum field theories, some of classical mechanics is briefly revisited, if only to re-establish terminology. Quantum field theory and its difficulties are then treated, giving several examples of why the study of cosmic strings (and therefore quantum field theory) through gravitational waves is incredibly important for general relativity.

This section is mostly based on [97–101]. For topology the reader can consult [102–104].

2.2.1 Classical Field Theory

In classical mechanics, the study of the mechanical motion of objects, there are two major and equivalent formalisms that reformulate the laws of Newton [105]. Hamiltonian mechanics describes the evolution of a system by a measure of the total energy called the *Hamiltonian*, defined as the sum of kinetic and potential energy. This formalism can be reformulated in terms of symplectic geometry, and by geometric quantization the classical formalism can be upgraded to a quantum formalism [106]. This is however foregone in this section. For Lagrangian mechanics, the role of the Hamiltonian is taken by the *Lagrangian*, given in contrast by the difference of kinetic and potential energy. Both flavours will be important in this section.

Considering the mechanics of a collection of n discrete objects such as particles, one can picture these objects as being placed on a uniform lattice that spans Minkowski spacetime. A field will then arise if the continuum limit $n \to \infty$ is taken within the confines of the lattice. The observable properties become continuous, where every $x \in \mathbb{R}^{3,1}$ can now be assigned a degree of freedom $\phi(x)$, and this ϕ is then exactly the classical *field*. In line with special relativity, such a field is called *relativistic* if it obeys the laws of special relativity. An important example of a relativistic field is given by classical electromagnetism [107], archetypical to the extent where the fields in classical field theory are named after the electrical and magnetic fields.

In this new setting of spacetime, a Lagrangian L can be upgraded to a *Lagrangian density*

$$\mathcal{L} = \int_{\mathbb{R}^3} L \, \mathrm{d}^3x \tag{2.41}$$

by integrating out the spatial dimensions. This density is then used as an ingredient to define the action (that describes how energy changes with trajectories) by also integrating out time, meaning the action itself is a functional as a definite integral of a Lagrangian density. For the purpose of this section, the Lagrangian density itself will be taken as the basis for quantum field theory, as Lagrangians are manifestly Lorentz-invariant, whereas Hamiltonians are not.

2.2.2 Quantum Field Theory

Quantum mechanics is a theory finding its origin in the wave-particle duality. Experimentally, it was found that even though classical mechanics assumes light to be either behaved like waves or particles, it in fact behaves like both [108]. This contradiction eventually lead to the development of a new theory in quantum mechanics, named after the mechanics of the *quanta*, the irreducible building blocks of nature. The duality is addressed in this theory by formulating the wave function, that can give the position of a particle and evolves in time according to a wave equation. This variation adds a level of uncertainty that manifests probabilistically, and it is this stochasticity that marks quantum theory as a departure from classical theory.

Quantum field theory then sprung from the need to combine quantum mechanics with relativistic fields in such a way that the resulting theory obeys the same transformational requirements as does special relativity. It is therefore a theory that marries the triple of special relativity, classical field theory, and quantum mechanics. Efforts have been made to incorporate curved spacetimes into quantum field theory, but have as of yet not yet seen success. Regardless, with both special relativity and classical field theory covered, quantum mechanics is up in the roll call. This theory will be introduced through five postulates.

Postulate 2.2.1 (Postulate 1)**.** The *state space* of a quantum mechanical system is given by a Hilbert space $\mathcal{H}$. A state is therefore a unit vector $\psi \in \mathcal{H}$, with states on the same projective ray representing the same physical state.

An example of a state can be the position of a single particle, specified as a unitary element of some predetermined Hilbert space. The generic inner product of this space will be denoted by $\langle \cdot, \cdot \rangle$.

Postulate 2.2.2 (Postulate 2)**.** For every classical and real-valued observable quantity f as a function on the classical phase space there exists an associated self-adjoint operator $\hat{f}$ on $\mathcal{H}$.

An example of such a function f can therefore be the position of a particle, but the examples will also include properties like the energy, momentum and angular momentum. The upgrading of functions to operators is called (first) *quantization*, with "first" for reasons that will become apparent later in this subsection. Considering the wave-like time evolution of the next postulate, first quantization can be interpreted as the assertion that particles behave like waves.

Postulate 2.2.3 (Postulate 3)**.** For the classical Hamiltonian H with associated operator $\hat{H}$, and the Planck constant $\hbar$, the time evolution of a state ψ is governed by the Schrödinger wave equation:

$$\frac{\mathrm{d}\psi}{\mathrm{d}t} = \frac{1}{i\hbar}\hat{H}\psi. \tag{2.42}$$

The solutions of this famous equation are called *wave functions*.

Postulate 2.2.4 (Postulate 4)**.** If a quantum system is in a state $\psi \in \mathcal{H}$, the probability distribution for the measurement of the classical observable f satisfies the following equality for its moments:

$$\mathbb{E}[f^m] = \langle \psi, (\hat{f})^m \psi \rangle. \tag{2.43}$$

This postulate establishes that the square modulus of ψ can be considered to be a probability density function, and the observable is integrated against this function to obtain the expected value. This invites probability and measure theory, and leads to the famous concept of waveform collapse. A measurement will lead to a single value, and at that exact point in time, the wave function is said to have collapsed onto this value. This is ascertained by the fifth and last postulate.

Postulate 2.2.5 (Postulate 5)**.** Assume that a quantum system resides in a state ψ, and that an observable f with associated operator $\hat{f}$ is measured. If the result of this measurement is $\lambda \in \mathbb{R}$, the system will then be in a state ψ' satisfying $\hat{f}\psi' = \lambda\psi'$ immediately after.

This postulate introduces the importance of eigenvalues. Quantum mechanics is typically performed in terms of eigenvalues and eigenvectors, relating to the discrete nature of the quanta. This is made concrete in for instance the study of the harmonic oscillator, where a ladder of energy levels emerges as an operator spectrum. This algebraic description relates to the ladder in the definition of the spin-weighted spherical harmonics in subsection 2.1.6.

Different interpretations of this probabilistic nature exist, and not one is currently preferred. Two examples are the realist interpretation, that argues the theory is incomplete, or the Copenhagen interpretation, that believes a measurement forces the system into a single state. The author would suspect incomplete information, as our local spacetime is only Minkowskian in approximation due to the presence of matter. Regardless of the interpretation, quantum theory is employed for many applications. Examples include the reduction of quantum noise in gravitational-wave detectors [34] or quantum computing [109].

Using these postulates a high-level description can immediately be given of how special relativity and quantum mechanics fail to combine. Evaluating the probabilities of particle positions in the exterior of the future light cone may result in non-zero probabilities [99], and there is no mechanism that prohibits this. Relativistically, this is ofcourse unacceptable. An alternative reasoning is that there exist quantized observables for position, but no operator exists for time. Relativity assumes that time and space stand on equal footing, meaning this creates a qualitative objection.

A more involved relativistic example is found through the information-theoretical study of black holes. Hawking radiation is a proposed thermal black-body radiation outside of the event horizon of a black hole, which machinates a gradual evaporation of the black hole [110]. Once a quantum state is inserted into a black hole, the information should therefore radiate thermally through this mechanism, although at random. The argument for this depends on the no-hair theorem. Because of this non-determinism, this process should be irreversible, or equivalently, non-unitary. This is, however, in direct violation of the third postulate of quantum mechanics. The theory of differential equations asserts that under the assumption of existence, the solution must be an exponential, and therefore the evolution of the state should evolve unitarily. This conundrum is called the black hole information paradox [111].

Although it will not remedy the latter paradox in curved spacetimes, the formulation of a quantum theory in terms of relativistic fields can combat the previously discussed faults. Second quantization, or canonical quantization, is the process of upgrading fields to operator-valued fields. This creates fields of quantum observables, or the quantum fields that are the subject of quantum field theory.

At the time of writing, there is no axiomatic definition of what a general quantum field theory is. A notable exception is the specific case of topological quantum field theory, defined category theoretically through the Atiyah-Segal axioms [112]. These axioms describe how compact oriented manifolds and their cobordisms are mapped onto Hilbert spaces and their bounded linear operators through a functor that preserves the monoidal and dagger structures. This formalism in fact gives rise to a mathematically formal method of quantum computing in topological quantum computing [113]. In this subsection, the Gårding-Wightman axioms are used as a guideline for mathematical rigour [114]. It is worthy of note that the difficulty of axiomatising quantum field theory can be demonstrated by the same line of reasoning as was used for solving the Einstein field equation, as the axioms are explicitly mentioned in the Yang-Mills Millennium Problem [39].

A *quantum field theory* will then be defined as a map:

$$\Phi : \mathcal{C}(M) \to \mathcal{B}(\mathcal{H}) \tag{2.44}$$

that sends smooth and compactly supported functions on a spacetime M to linear operators on a Hilbert space $\mathcal{H}$. In more detailed treatments, this map is introduced as an *operator-valued distribution*, a distribution in the same way the famous Dirac delta-distribution is. As will be treated shortly, an example of a quantum field operator can be the creation operator $\phi_x^\dagger$ that, when applied to the vacuum state, "creates" a particle at a given point x. This is problematic however, as a particle can never truly be localised to a point. Instead, the particle must be assigned to a region of space. This assignment addresses an issue referred to as ultraviolet divergence, named after the ultraviolet catastrophe [115]. In a way, the operator has to be assigned a volume V for the particle to reside in. This is done by a process called smearing, where the operator is integrated against a test function, known from the theory of distributions [116]. As such, for $\phi_x^\dagger$, the resultant operator for a test function $f \in \mathcal{C}(M)$ should take the form

$$\int_V f(x)\phi_x^\dagger \mathrm{d}x \tag{2.45}$$

for some volume V. An example of f can be a Gaussian. In lesser generality this can be expressed as an operator field on Minkowski space by a map:

$$\Phi : \mathbb{R}^{3,1} \to \mathcal{B}(\mathcal{H}). \tag{2.46}$$

It is theorised that every type of particle admits its own type of quantum field as part of the *standard model* [99,100], and that particles manifest themselves as excitations of these fields. Second quantization can therefore also be interpreted as establishing that waves behave like particles, or the reverse implication of first quantization. In order to formalise these statements the Fock space is defined. The Fock space allows for the use of the occupation number representation, where states are described by the number of particles present.

Before defining the Fock space, two types of particles are distinguished based on their commutation relations. The *bosons* are those particles that, when exchanged in position, leave their wavefunction unchanged. For *fermions* such an exchange introduces a parity transformation to the wave function, meaning a change of sign. The definition of the Fock space accounts for these differing statistics, and this difference manifests itself in the commutation relations of the operators defined on this space soon after.

Definition 2.2.1 (Fock space)**.** For $p \in \{+, -\}$, a Fock space $F_p(\mathcal{H})$ based on a Hilbert space $\mathcal{H}$ is the Hilbert space completion of:

$$\bigoplus_{n=0}^{\infty} S_p \left[\mathcal{H}^{\otimes n}\right], \tag{2.47}$$

where S_+ is an operator that symmetrises a tensor product to account for bosonic statistics, and S_- is an operator that anti-symmetrises a tensor product for fermionic statistics.

If the Hilbert space is taken to be the quantum state space of a single particle, then the Fock space is the extension to the many-particle state space with particles of the same type. The n-th component of the direct sum in the above definition is referred to as the n-th particle sector of the Fock space, and refers to a state wherein n particles exist. One can shift to and from these sectors by defining creation and annihilation operators on the Fock space (that generate a full set of basis states), and this is in fact how the state of a quantum field is determined. The map given in Eq. 2.46 yields an operator that can be expanded in terms of the aforementioned operators, and when applied to the vacuum 0-sector gives the state of the system. The quantum field is said to be excited at a point in the spacetime when the application of the operator gives a state residing in a non-zero k-sector of the Fock space. This corresponds to the existence of one or multiple particles at that point.

As has been illustrated before, the starting point for new physical theories have often been the theories that have come before. So too for quantum field theory, where canonical quantization requires a classical field to have already been constructed. Although the machinery itself will not be discussed, it is important to remark that this starting point is taken to be that of a Lagrangian density. The reason for this is its relativistic invariance, as well as its connection to Noether's theorem [99]. The use of the Lagrangian density therefore carries over to the study of quantum fields.

In the wake of wave-particle duality, one could question the implications for gravitational waves. An answer to this question would need to establish that quantum theory is applicable to the regime of gravitational waves, or in other words, that gravitational fields are quantum fields. If so, string theory could predict the *graviton* as the dual particle for gravitational waves [117], bearing the same relation as photons and electromagnetic waves. String theory is a quantum theory that essentially replaces localised points by strings, and in the limit reproduces general relativity [117, 118].

As a closing remark on quantum field theory, a correspondence is conjectured between the Anti-de Sitter spaces mentioned in subsection 2.1.1 and conformal field theories, those quantum field theories that are invariant under locally angle-preserving maps. This allows for similar constructions to the fourth postulate, where a classical observable is computed through an operator-correspondence [119].

2.2.3 Cosmic Strings

Intuitively, internal symmetries are those symmetries that leave a system invariant from an outside perspective. In the case of a quantum field, with G a group, the field is said to have an *internal* and *global* G-symmetry if the Lagrangian density describing the field is invariant under the left action of G, so that for all $g \in G$:

$$\mathcal{L}(g \cdot \phi) = \mathcal{L}(\phi). \tag{2.48}$$

The local counterpart to these global symmetries are the *local symmetries*, where the action of $g(x)$ is parametrised by elements of the spacetime. Naturally, a Lagrangian is likely not to be invariant under local transformations. Analogously to the gauge freedom mentioned in the coverage of general relativity, one can add a *gauge field* that counteracts the local group action. The geometric interpretation is that this field can be absorbed into the covariant derivative, therefore changing the connection on the manifold.

Consider next the zero-set or vacua of the Lagrangian density, which is assumed to be a smooth manifold, and is therefore referred to as the *vacuum manifold*. In terms of global symmetry it is possible that the vacuum manifold, as a submanifold of the spacetime, admits a lower degree of symmetry than does the spacetime. If this is the case, and the lower symmetry is given by a subgroup $H \subseteq G$, the G-symmetry is broken down to an H-symmetry once the system assumes a minimising state ξ. The unbroken symmetry group H will then be the stabiliser subgroup of the state ξ in G, and this process is called *symmetry breaking* [98]. A related and relevant result is Goldstone's theorem [98], that describes the bosonic particles that arise following symmetry breaking.

The breaking of symmetry can occur as the result of a phase transition [101, 120]. As an example, if a ferromagnetic material is cooled past its Curie point, the system undergoes a phase transition. After this transition, the material will show different regions of magnetisation that are separated by domain walls, experimentally observed in condensed matter [99, 121]. Domain walls are examples of topological defects, or topological solitions, as they can be traced through topology [121]. This is explained in generality in the next paragraph.

A topological defect is introduced if during the breaking of symmetry the homotopy groups of the vacuum states have changed, or in other words, if at least one of the homotopy groups of the orbit $H \cdot \xi$ differs from the corresponding group for $G \cdot \xi$. In generality this can be phrased by stating a *topological defect* occurs in a theory if there exists a solution to a system of partial differential equations that is homotopically distinct from the vacuum solution. In terms of physical intuition, if $\pi_k(G \cdot \xi)$ was assumed trivial, the breaking of symmetry results in the removal of vacuum states that are detected by null-homotopic S^k-embeddings. Due to the Lagrangian density describing energy, a concentration of energy must have been introduced to the system at these positions, and it is these concentrations that are labeled topological defects.

Cosmic strings are hypothetical one-dimensional topological defects, registered in the fundamental group π_1. The explanation of their name comes in two-fold. First, these defects are cosmological, as their appearance may have been caused by the phase transitions associated with grand unification [120]. This directly motivates cosmic string searches, as they can be considered a vehicle for learning about the early universe. Second, in order for the contraction of an S^1-embedding to a base point to fail, the obstructing defect must be one-dimensional, and therefore string-like.

There exists a parallel string-theoretical class of cosmic strings [120]. The possibility of the elemental strings blowing up to cosmological scales has been theorised, which would lead to *cosmic superstrings* that should behave very similar to the cosmic strings introduced here. One specific difference is found in their interactions, where the intercommuting probabilities upon string interaction differs. These probabilities will aproach unity for cosmic strings, whereas for cosmic superstrings this probability will depend on the string parameters [122].

Physically both types of cosmic strings would manifest as extremely thin but massively dense strings. Because of the assumption of zero-thickness their large-scale dynamics can be described by the Nambu-Goto string action [117, 120], and these assumptions have been used to simulate the evolution of large networks of cosmic strings [123–127]. Most importantly for this thesis, their density should enable gravitational effects. One such effect is caused by an angular deficit in the local geometry [101], resulting in gravitational lensing of background light sources. Gravitational lensing here means that cosmic strings would alter nearby light geodesics, making the cosmic string act as a lens between an observer and the light source [10]. Due to the string action, there are also dynamical mechanisms for gravitational-wave generation on the string. These are string cusps, string kinks, and kink-kink collisisons [128]. This thesis will specifically rely on cusps.

A *cusp* is a point on a curve where the derivative vanishes, with left limit ∞ and right limit $-\infty$ at this point. The tangent vector will therefore vanish and change sign at the cusp. On a relativistic string this would see the string snap into a cusp shape, instantaneously accelerating the point to the speed of light and emitting a burst gravitational-wave in the direction of acceleration. The waveform of this event on a closed loop of length l and string tension μ has been phenomenologically computed as a powerlaw function of the frequency f as [128]:

$$\begin{aligned} h_{l,z,\mu}(f) &= \left[(2/3)^{2/3} 8/\Gamma^2(1/3) \frac{l^{2/3}\mu}{(1+z)^{1/3} r(z)}\right] f^{-4/3} \\ &\approx \left[\frac{0.85 \cdot l^{2/3} \cdot \mu}{(1+z)^{1/3} r(z)}\right] f^{-4/3}. \end{aligned} \tag{2.49}$$

In this equation z is the redshift accounting for the stretching of wavelengths due to the relative motion of the cosmic string and observer, $r(z)$ is the co-moving distance between the two, and Γ denotes the Gamma function. Waveforms for kinks follow a similar powerlaw, with power $-5/3$ for individual kinks and -2 for collisions [128].

Chapter 3

Data Analysis

The objective of this chapter is the building of a high-level overview of the state of gravitational-wave data analysis. The data analysis concepts that underly the following chapters in this thesis are therefore introduced here, starting with signal processing and the software that implements these techniques, before moving to general machine learning and statistics.

3.1 Signal Processing for Gravitational Waves

This section collects topics from signal processing that are applicable to gravitational-wave data analysis, and the concepts are therefore directly applied in this context where appropriate. The subsections leads up to two main lines of research: *searches*, that are looking to recover signals from data, and *parameter estimation*, the process of estimating the parameters that generated recovered signals. The treatment presented in this subsection is by no means complete and only samples the required concepts.

3.1.1 Time and Frequency Domains

Thus far, the Fourier transform has been used simply as an isomorphism between functions in the time domain and frequency domain. However, when analysing data, it is useful to make an informed choice of domain. For this reason both options are briefly reviewed here. An example of a waveform is shown in both the time and frequency domains in Fig. 3.1.

Many operations are easier in the frequency domain. An example can be seen in ordinary differential equations, where initial value problems can be mapped onto algebraically solvable polynomial equations using the Laplace transform. More generally, the convolution theorem maps convolutions in the time domain onto pointwise multiplication in the frequency domain. In addition, the frequency domain might reveal cyclic behaviour in a signal that is not apparent in the time domain, and allows for the splitting of a signal into the amplitude and phase.

A major advantage of the time domain however is that the direct time evolution is easier to interpret directly, and as a consequence, it is easier to separate a signal

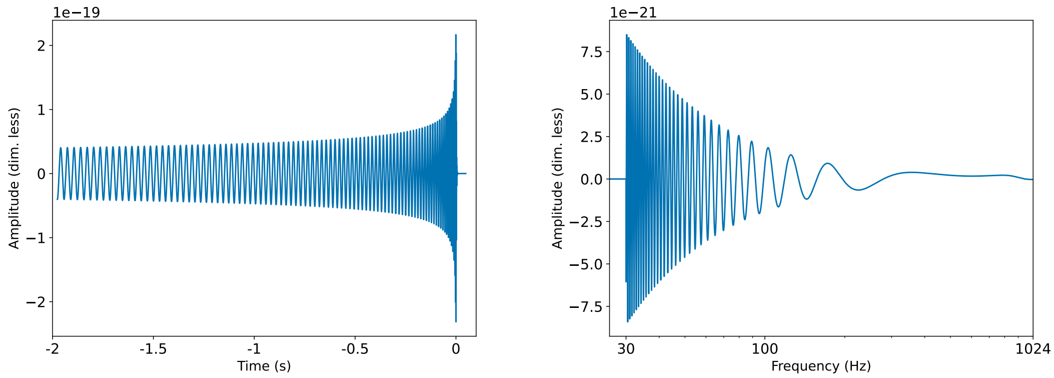

(a) A waveform in the time domain. (b) A waveform in the frequency domain.

Figure 3.1: An example of the plus-polarisation of a waveform generated in both the time and frequency domains for a generic system with $m_1 = m_2 = 10M_\odot$ and a lower frequency cut-off of 30 Hz. The time-domain waveform on the left was generated using the `IMRPhenomT` approximant, and the frequency-domain waveform on the right was generated using the `IMRPhenomXAS` approximant.

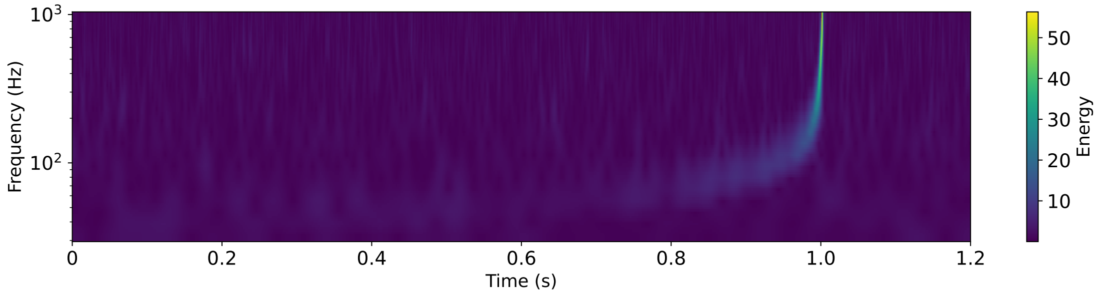

Figure 3.2: The spectrogram of the frequency-domain waveform shown in Fig. 3.1b injected into Gaussian white noise before the strain was whitened (these concepts are treated in Subsec. 3.1.2). This figure shows the chirping increase in pitch of a compact binary coalescence waveform.

from noise by eye. It is also costly to apply variations of the Fourier transform to map data into the frequency domain, even when for instance the Fast Fourier Transform is used [129]. This is a major consideration for optimisation as strains are collected in the time domain.

It is to an extent possible to combine the two domains in time-frequency representations. The Q-transform results in an example of such a representation [130]. Whereas the discrete Fourier transform will assume the use of equal spacing, the Q-transform uses geometric spacing. This means that the spacing will be equal in a logarithmic basis, and the transform can be used to show the evolution of the frequency spectrum through time as a heatmap. This visualisation is called a *spectrogram*, and is shown for the waveform in Fig. 3.1b in Fig. 3.2. As the heatmap represents the energy in the form of the amplitude, information on the phase is mostly lost.

In time-frequency analysis, and more generally in signal processing, there is a degree of uncertainty referred to as the Gabor limit [131]. Interestingly, this mirrors the Heisenberg uncertainty principle that concerns the standard deviations of position and momentum from quantum mechanics [97].

3.1.2 Noise and Signal Recovery

In this subsection $X = \{X_t\}$ will denote a complex continuous stochastic process with time indexed by $\mathbb{R}$. The results derived in this subsection will translate to the discrete case indexed by $\mathbb{Z}$ through a change of measure. A signal will be a measurement or realisation of the stochastic process.

The process X is called *weak-sense stationary*, as opposed to strictly stationary, if both the mean and covariance of the process have no time-dependency, and moreover, the second moment $\mathbb{E}[X_t^2]$ is finite for all t. If this is the case there must exist $\mu, \sigma \in \mathbb{R}$ so that for all t:

$$\mathbb{E}[X_t] = \mu \text{ and } \mathrm{Var}(X_t) = \mathbb{E}[X_t^2] - \mathbb{E}[X_t]^2 = \sigma^2. \tag{3.1}$$

These properties imply that for a weak-sense stationary process the *autocovariance function*, with the asterisk denoting complex conjugation,

$$\mathrm{K}_{XX}(t, t+\tau) = \mathrm{Cov}[X_t, X_{t+\tau}] = \mathbb{E}[X_t X_{t+\tau}^*] - \mu\mu^* \tag{3.2}$$

will be an even function depending only on the time difference (expressed here as the time lag τ), and not the individual times. This can be verified by writing out the definition of the covariance and making use of the linearity of the expectation. Adding the term $\mu\mu^*$ then yields the *autocorrelation function*

$$\mathrm{R}_{XX}(\tau) = \mathrm{R}_{XX}(t, t+\tau) = \mathrm{K}_{XX}(t, t+\tau) + \mu\mu^*. \tag{3.3}$$

For general time series, and therefore in particular stochastic processes, the autocorrelation function gives the correlation of a time series with a delayed copy of itself. It can therefore be useful to describe noise, as the memory component introduced by the time lag introduces a dependency between different sampling times. In particular, the *power spectral density* S_X is defined as the Fourier transform of the autocorrelation function:

$$\mathrm{S}_X(f) = \tilde{\mathrm{R}}_{XX}(\tau) = \int_{-\infty}^{\infty} \mathrm{R}_{XX}(\tau) e^{-2\tau\pi i} \mathrm{d}\tau, \tag{3.4}$$

and is also commonly written $\mathrm{PSD}(f)$. Intuitively speaking, the value of the power spectral density at a frequency f describes the average power contained within the frequency of the signal. It is an even function, inheriting this property from the autocovariance, and the literature therefore often uses the one-sided power spectral density.

In order to more clearly explain why it is called the power spectral density, and in fact describes power, the general energy $\mathrm{E}(s)$ contained in a time series s is defined as the integrated squared modulus:

$$\mathrm{E}(s) = \int_{-\infty}^{\infty} |s(t)|^2 \mathrm{d}t = \int_{-\infty}^{\infty} |\tilde{s}(f)|^2 \mathrm{d}f \tag{3.5}$$

where the second equality follows from Parseval's theorem, or more generally the Plancherel theorem. This motivates the definition of the energy of the stochastic process to be given by

$$\mathbb{E}[X_t^2] = \mathrm{R}_{XX}(\tau = 0) = \int_{-\infty}^{\infty} \mathrm{S}_X(f) e^{2\cdot 0 \cdot \pi i} \mathrm{d}f = \int_{-\infty}^{\infty} \mathrm{S}_X(f) \mathrm{d}f, \tag{3.6}$$

or the inverse Fourier transform at $\tau = 0$. Note that the domain of integration can be limited, and the integral will therefore then give the power at a range of frequencies, with $\mathrm{S}_X(f)$ denoting the power at an isolated frequency f.

Alternative definitions of the power spectral density exist, for instance the definition given in [34]. Such definitions are typically equivalent following application of the Wiener-Khinchin theorem [132].

In gravitational waves the graph of the power spectral density for a detector is commonly called its *sensitivity curve*, as the power spectral density holds information on the noise floor. As the power spectral density is given in units of Hz^{-1}, a signal for which the graph of its waveform lies at least partially above the sensitivity curve is conceptually detectable. This however does not exclude signals with graphs below the sensitivity curve, which can still potentially be detectable. Studies of detector sensitivities are a part of detector characterisation [133].

In addition, the power spectral density is used to colour noise and data, with different colourings representing different properties of the resultant strains. In the case of white, the noise is characterised by having a flat frequency spectrum, or having equal power in every frequency in the spectrum. If the power spectral density of a detector is therefore set to be fixed, and a strain is coloured to white by the power spectral density, one would expect the spectrum to be flat. However, if a signal is present in the noise, this will not be flattened, and stand out in the result. The process of dividing data in the frequency domain by the the power spectral density is therefore called *whitening* and is extremely helpful in aiding the process of finding signals in noise.

Injection of signals are also typically done in whitened data. After whitening, in the time domain, the injection of a signal amounts to the time series summation. This can be used to generate datasets when sets of sufficient size are not on hand, by generating noise following a power spectral density and injecting artificially generated signals.

The discussion of whitening leads into the process of *matched filtering* and the *signal-to-noise ratio*. Matched filtering is a method from signal processing that is known to be optimal for the recovery of a given filter from Gaussian and stationary noise [134]. In order to achieve this, matched filtering convolves the filter with the data to obtain a scalar quantity called the signal-to-noise ratio. For a function space into which the filters, or waveforms, can be embedded, the ratio is computed through the Hermitian inner product:

$$\langle u, v \rangle = 4\mathrm{Re} \int_0^\infty \frac{u(f)v(f)^*}{\mathrm{PSD}(f)} \mathrm{d}f. \tag{3.7}$$

The intuition is that the numerator gives the dot product as a quantity giving the correlation of the waveform and data, and the denominator then normalises this frequency content. This quantity can therefore be seen as giving the presence of u in v, as a ratio of signal to noise. One should note that for some applications the quantity is further normalised by dividing over the time duration of the waveform, a factor $\langle u, u \rangle$, or both.

Typically the waveform one is looking for is not known. Instead, multiple *templates* corresponding to system parameters can be used, and these templates are collected in what is called a *template bank*. Such a template bank can for instance be generated by specifying sets of parameters and an approximant [135]. The name of matched filtering can then be dissected into matching the best filter to the data.

3.1.3 Searches

Gravitational-wave searches will follow typical workflows that involve both signal processing and statistical techniques to detect and verify transients in data. In this subsection this process is briefly reviewed.

A distinction is made between two types of searches. Modelled searches are those searches that assume signal models are available, and unmodelled searches are those that do not make this assumption. Examples of pipelines for both will be mentioned in the next subsection on gravitational-wave software. In this subsection the focus is on modelled searches based on matched filtering. These searches can be performed online on incoming data, or offline on historic data.

In searches, the order of magnitude of the number of templates n required for a procedure is often denoted by $\mathcal{O}(c)$ or $O(c)$ for $c \in \mathbb{R}$, which is to be read as $n \propto c$. This notation is used interchangeably with the Bachmann-Landau symbol for asymptotic domination and algorithm complexity classes [136], defined for $g : \mathbb{R} \to \mathbb{R}$ by:

$$\mathcal{O}(g) = \{f : \mathbb{R} \to \mathbb{C} \mid \exists\, k, x_0 \in \mathbb{R} \text{ s.t. } x \geq x_0 \implies |f(x)| \leq kg(x)\}. \tag{3.8}$$

As it can be proven from this definition that for all $c' \in \mathbb{R}$ it holds that $\mathcal{O}(g) = \mathcal{O}(c'g)$, and hence for all $c_1, c_2 \in \mathbb{R}$ it must hold that $\mathcal{O}(c_1) = \mathcal{O}(c_2)$, the reader should take care to tell the difference in notation usage from the context. In this thesis the definition given in Eq. 3.8 will be upheld consistently.

A matched-filtering-based search pipeline will produce statistics such as the signal-to-noise ratio or the χ^2-statistic [137]. If these statistics attain sufficiently high values, a *trigger* is produced, and at this time the template and corresponding GPS timestamp are saved. These statistics can be used to form ranking statistics, and these will form the basis for deciding coincidence when combined over different detectors in the network. Following this, detection statistics can be computed that can verify a result as a *gravitational-wave candidate*. Finally, combined with tresholds on the detection statistics one can speak of a *gravitational-wave event*, and the follow-up procedures will typically include data quality research to verify data standards, and parameter estimation, where the parameters for an event are estimated to a higher precision than a coarse search can determine.

3.1.4 Parameter Estimation

The result of a search will, from the previous subsection, typically be a waveform template that was matched to detector data. Template banks are typically generated from coarse samplings over the parameter space, and the parameters that generated the triggering template will therefore not give the best possible estimation of the true parameters. Instead, parameter estimation is designed for this purpose, finding its origin in Bayesian statistics.

Let A, B denote two events with probability measure $\mathbb{P}$ that absorbs the assumption of a model hypothesis. Recall then that Bayes' theorem poses the equality

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(B|A)\mathbb{P}(A)}{\mathbb{P}(B)}. \tag{3.9}$$

Under Bayesian inference the probabilities in this equation will have interpretations that express the beliefs that certain events will come to pass. In particular, the conditional probability $\mathbb{P}(B|A)$ is the likelihood of B occuring following the event A. If one sets B to signify a strain measurement of h, and A to θ, short for a system with $\theta \in \mathcal{P}$ generating a signal s present in $h = n + s$, this quantity can be optimised in θ. Note that the hidden hypothesis then encodes the assumptions put on the system and waveform model. If it is assumed the noise n is Gaussian and s is not identically zero, the likelihood obeys

$$\mathcal{L}(h, \theta) = \mathbb{P}(h|\theta) \propto \exp\left[-\frac{1}{2}\langle h - w(\theta), h - w(\theta)\rangle\right] \tag{3.10}$$

in the frequency domain [138], where the inner product is given in Eq. 3.7, and w

is the waveform model. The likelihood therefore stands in strong relation to the signal-to-noise ratio. Furthermore, as the natural logarithm is a strictly increasing function, the maximisation of the likelihood is equivalent to the maximisation of the log-likelihood $\ln \mathcal{L}(h, \theta)$. This reduces the need for exponentiation, reducing computational cost.

The current parameter estimation workflow is therefore to establish a sampler over the parameter space that maximises the likelihood function by repeatedly sampling sets of parameters, computing the corresponding waveforms, and computing the likelihood until an extremum is found. The most costly step of this routine is evidently the computation of waveforms.

In addition to the likelihood, the priors and posteriors derived from Bayes' theorem take in a central role in the verification aspects of parameter estimation. The reader is referred to [138, 139] for more thorough treatments that includes these components.

3.2 Gravitational-Wave Analysis Software

This very brief section covers common software used in the field of gravitational waves, covering general packages, search pipelines and other libraries. This listing is not exhaustive, and moreover, software packages are always in development. Several software suites overlap in their applications. If this is the case, the software is categorised in the subsection thought to be most applicable.

3.2.1 General Packages

The `GWOSC` package, short for Gravitational-Wave Open Science Center [140, 141], is a `Python` interface to call on the open data releases of the LVK observatories. This includes the `GWTC` (Gravitational-Wave Transient Catalog) that catalogues the transients found in observatory data [42, 142–144].

A major package used as the basis for many other packages is `LALSuite` [145], short for the LIGO Scientific Collaboration Algorithm Library Suite. This package was written in `C` but also comes wrapped in `Python`, and contains many functions and routines in separated namespaces. Examples include `LALSimulation` for simulations, or `LALBurst` for burst gravitational-wave analysis.

Connected to `LALSuite` is `GWpy` [146]. This package was designed to simplify more extensive packages such as `LALSuite` and includes many examples.

Contracting the terms `Python` and compact binary coalescence, `PyCBC` [147] is a collection of algorithms designed for searches, parameter estimation, and other forms of processing of signals from compact binary coalescences. It is however also useable for other types of sources, and was extended to a search pipeline [148].

3.2.2 Search Pipelines

`GstLAL` [149] is a combination of `LALSuite` and `Gstreamer`, a `Python` multimedia framework [150]. As gravitational waves and multimedia are both applications of the more general theory of signal processing, `Gstreamer` can be used for the streaming and processing of gravitational waves, even though it was designed for multimedia applications. `GstLAL` is then a package that connects `LALSuite` to the functionality provided by `Gstreamer`, used for the design of matched filtering search workflows for both compact binary coalescence and burst searches.

`BayesWave` refers to both an algorithm [151] and an unmodelled search pipeline [152] based on this algorithm that uses Bayesian statistics to find transient signals and separate them from glitches. In order to do so `BayesWave` uses wavelets [153] to reconstruct and study the transients. Wavelets are briefly oscillating signals that are useful in the face of the Gibbs phenomenon.

Using detector network coherency in search for burst gravitational waves, `coherent WaveBurst` is an unmodelled search pipeline for both the detection and reconstruction of signals [154]. This pipeline also uses wavelets, but does so in time-frequency analyses.

3.2.3 Parameter Estimation

Several other packages such as `PyCBC` can be used for parameter estimation, but `Bilby` [155] was designed with the sole objective of performing parameter estimation. It does so based on Bayesian inference, and was initially meant only for the parameter estimation of compact binary coalescences. However, it is also applicable to the parameter estimation of burst gravitational waves.

3.3 Machine Learning and Statistics

This section will introduce the core concepts from machine learning and statistics used in the following chapters of this thesis. At a certain point in time machine learning was seen as a subfield of statistics, called statistical learning. One can however argue that if there is still a difference, the fields have diverged in their approach. The material in this section is based on [156–160].

3.3.1 Neural Networks

Neural networks are inspired by the way science understands the human brain to work [160]. One of the fundamental units of the brain is the *neuron*, an excitable cell that is subdivided into the soma containing the nucleus, and two types of connecting nerve fibers that weave the neuron into a larger network of neurons [158]. The first connecting fiber is the *dendrite* that carries information to the soma. The second type of fiber is the *axon*, that carries information out of the soma. From this description one can see how a network of neurons might lend itself to the use of being a computational unit.

The first step towards a realisation of such a unit was the McCulloch-Pitts model of a single neuron [161]. Mathematically, the model is a function $m : \{0,1\}^n \to \{0,1\}$ that aggregates a discrete number of binary inputs representing the dendrite throughput into one tresholded binary output or axon:

$$m(x_1, ..., x_n) = I_\alpha \left(\sum_{i=1}^{n} x_i \right), \tag{3.11}$$

where $I_\alpha(x)$ is the indicator function that is 1 if and only if $\alpha \leq x$, and 0 otherwise.

It was later discovered that the bonds between neurons that fire at the same time strengthen. Following this biological advancement, the model was extended to the Rosenblatt perceptron [162]. For the perceptron, the input domain is no longer limited to binary values, a bias is added, and the inputs are weighted:

$$p(x_1, ..., x_n) = I_\alpha \left(b + \sum_{i=1}^{n} w_i x_i \right). \tag{3.12}$$

The perceptron, with general *activation function* γ instead of I_α, is the basis for the modern neural network, and therefore referred to as an *artificial neuron*, or just a neuron. These neurons are then woven together into a network of neurons to form the *neural network*.

It is sometimes useful to interpret neural networks as graphs [163], where the vertices represent the soma of the neurons, and the directed edges represent the fibers. The terminology of graphs will therefore be used for neural networks in this subsection. A *feedforward neural network* is then a directed acyclic graph, and the neurons can be partitioned into sets called layers based on their in-degree and out-degree. The input layer consists of those neurons with an in-degree of zero and an out-degree of at least one. On the other extreme, the output layer contains neurons with a non-zero in-degree and zero out-degree. The internal or hidden layers contain all internal neurons with both an in-degree and out-degree of at least one, with the assigned layer corresponding to the shortest path length from a neuron in the input layer. A neural network is called a *deep neural network* if it has at least one single internal layer.

The versatility of neural networks is that they can approximate a large range of functions by fitting values to the neuron biases and weights. In particular, the inclusion of activation functions, like for instance the sigmoid function, allows a neural network to learn non-linearities. The process of fitting weights for a given task is called the *training* of the model. However, even for relatively shallow networks, the training is incredibly costly, as the parameter space is typically very high-dimensional. In terms of complexity it has been shown the training of a neural network is NP-complete [164]. Many optimisation algorithms therefore work approximately and iteratively. If f denotes the function that is to be learned, and θ_i collects all parameters in the i-th iteration, let $f_i := f_{\theta_i}$ denote the function defined by these parameters. The target is then that $\lim_{i\to\infty} f_i = f$.

In order to treat the training of a neural network as an optimisation problem, a cost function C needs to be defined. This function can also be called the loss function and will need to be optimised in the parameters. In the simplest case the cost function will have access to the true value $f(x)$ for a single datapoint x, and the optimisation problem is then the problem of finding:

$$\Theta = \arg\min_{\theta} C\left[f(x), f_\theta(x)\right]. \tag{3.13}$$

Typically this function will turn out to be non-convex, especially if multiple datapoints are included. This puts the problem of finding minima outside of convex optimisation where one can typically trust polynomial-time solution algorithms to exist. It is however well known in analysis that if the expression above is differentiable, it can be optimised iteratively by the gradient descent algorithm [159]. This algorithm repeatedly calculates and substracts the gradient, eventually stablising in a local minimum where θ_i and therefore f_i is determined. Note that this process is controlled by learning and decay rates that scale the gradient [159]. In practice stochastic gradient descent as well as backpropagation for the computation of the gradients are used [159], along with optimisers [165, 166].

The above paragraph assumed true values of f to be available during training. If this is the case, this is called a *supervised* learning procedure, in contrast to *semi-supervised* and *unsupervised* learning procedures [159]. If f is a discrete function representing data classes, $f(x)$ is called the label of x. In the supervised case where labeled data is provided this dataset is often split into the *training*, *validation* and *test* sets. The training set is the set of datapoints that is used exclusively for the training, or the optimisation of the parameters. The model is exposed to the data in batches during the training phase, and if all training data has been processed, an *epoch* concludes. The validation set is then used for the independent validation of the performance of the trained models, and finally the held-out test set is used to give an independent measure of performance. Requirements on the datasets for specifically desired results are formalised in PAC-learning [167].

The effectivity of neural networks can be explained by the universal approximation theorem [168–170]. Let $\Sigma(\gamma)$ denote the span of a neural network with activation functions γ and a single hidden layer, representing all functions this neural network can represent. If X is a topological space, the neural network is called a *universal approximator* on X if $\Sigma(\gamma)$ is dense in $\mathcal{C}^o(X, \mathbb{R})$. The universal approximation theorem then states that if the rectified linear unit or sigmoid function is taken for γ, the neural network will be an universal approximator on $[0, 1]^n$. This result stands in close relation to the Stone-Weierstrass theorem [171], which states that under the right conditions, functions on compact Hausdorff spaces can be approximated up to arbitrary precision by polynomial functions with uniform convergence. The universal approximation theorem as stated here only applies to a limited number of neural network architectures. In practice however, many variations exist, and so too there exist many versions of the universal approximation theorem [168–170, 172, 173].

Besides the standard feedforward neural networks explained above, there are currently two other major types of neural network architectures. One is the *convolutional neural network* [160], where the arithmetic in Eq. 3.12 is replaced by convolutions with filters. Such filters can either be designed or learned during training, and were conceived as representing features in visual data. Convolutional neural networks can include skip connections that connect different layers, or residual connections that combine the information of different layers. These terms are sometimes used interchangeably. The second type is that of the *recurrent neural network* [160], where cycles may occur. Such networks allow for a memory in the network and are typically used for sequential data.

Different libraries implementing neural network methods exist. Examples include `PyTorch` [174] and `TensorFlow` [175]. Both these libraries include `CUDA`, a parallel computing platform and interface by Nvidia [176]. Although operations were traditionally executed on the central processing unit, which is effective for serial processing, this platform performs operations on the graphics processing unit. This processing unit is well-optimised for parallel processing because of its traditional main use in computer graphics, and can therefore handle linear algebra particularly well. Use of this platform can therefore highly accelerate inference [177].

As a final note on neural networks, due to their large complexities and opaque training procedures, they are notoriously difficult to interpret, and they are often referred to as *black-box* models [178]. At the same time however, they have proven to be extremely efficient at solving many complex real world problems. It is therefore worthwhile to work on the interpretability of neural networks.

3.3.2 Evolutionary Algorithms

Evolutionary algorithms [158, 179] are a class of heuristic search methods based on the Darwinian principle of natural selection. In this model of biological evolution, the selection mechanism leads populations to adapt to their environment, and this philosophy inspires the use of evolution-based schemes for optimisation problems that are too difficult to solve directly. As in genetics, the genome data is represented internally by genotypes that contain the hereditary information of an organism, whereas the phenotypes describe the externally observable characteristics. In practice this means that candidate solutions are compressed and encoded in *genotypes*, and the full candidate solution *phenotype* can be constructed from the genotype. An evolutionary algorithm then seeks to steer the evolution of the candidate solutions to an optimum through iterative generations. Besides an injective mapping from the set of genotypes G to the set of phenotypes P, an evolutionary algorithm is governed by the functions given in the following paragraph.

First, a heuristic [136, 158] is required, and this takes the shape of an objective function $O : G \to \mathbb{R}$ that assigns each genotype a fitness score that measures its performance in the environment. In every generation the genotypes will need to be evaluated by this function in order to select the most promising candidates that will form the following generation. In order to do so, parents are selected through a parent selection function S that randomly draws two parents from the set of genotypes equipped with a chosen probability distribution. These parents then form offspring through a recombination function $R : G \times G \to G$. Note that this recombination function can include stochastic mutations of the genotype, which are useful to add small perturbations that can nudge a solution away from possibly local but non-global optima.

In order to start the evolutionary algorithm, a first generation of solutions needs to be instantiated. Once this set is realised, the functions O and $R \circ S$ can be used to iteratively create new generations, and this is repeated until a set of stopping conditions has been met. Examples of such conditions include the reaching of a certain generation, or the stalling of variation within subsequent generations [158].

In general, evolutionary algorithms offer methods of exploring solution spaces where the solutions themselves can be complicated objects. Given the right assumptions, subsequent generations are expected to improve on their predecessors as lesser performing solutions are removed from the gene pool, and this leads to well-performing solutions in practice. However, it is not guaranteed that an evolutionary algorithm will find a global optimum, and there are few theoretical results that explain the effectivity and convergence properties of evolutionary algorithms. One notable theorem that is often cited is the Schema theorem [180], that states that evolutionary algorithms will drive up the average fitness with each passing generation.

3.3.3 Principal Component Analysis

Principal component analysis is a dimensionality reduction method that linearly maps vectorial data onto a lower-dimensional subspace. This subspace is spanned by an ordered basis of vectors called principal components. Both the method and these components are introduced here in the style of [157].

Let $\{x_i\}_{1 \leq i \leq m}$ with $x_i \in \mathbb{R}^n$ denote a dataset of size m, giving rise to a data matrix $X \in \mathrm{M}(n, \mathbb{R})$ of row rank m with the data vectors as its rows. For $p \leq n$, a rank p approximation to X can be given by

$$\mu + V_p X \tag{3.14}$$

where $\mu \in \mathbb{R}^n$, and $V_p \in \mathrm{O}(n, \mathbb{R})$ is an orthogonal matrix of rank p. Approximations specified by (μ, V_p) can be found using the method of least squares on the reconstruction error. If this is done for $p = n$, the transpose of V_n will coincide with the right singular vectors in the singular value decomposition of X:

$$X = UDV_n^{\mathrm{T}}, \tag{3.15}$$

where $U \in \mathrm{U}(n, \mathbb{R})$ are the left singular vectors, and $D \in \mathrm{M}(n, \mathbb{R})$ is the diagonal matrix of singular values. The column vectors in $XV_n = UD$ are then the ordered *principal components*. For lower rank approximations, the matrix V_p in Eq. 3.14 can be obtained from V_n by setting the column vectors past index p to identically zero.

The right singular vectors are in fact eigenvectors of the normalised covariance matrix of X, and as an alternative derivation after centering the data, this matrix can be decomposed in the eigenspace as

$$(m-1)^{-1}X^{\mathrm{T}}X = V_nLV_n^{\mathrm{T}}, \tag{3.16}$$

where $L = (m-1)^{-1}D^2$ is the diagonal matrix of decreasing eigenvalues. This follows from substitution of the expression for X given in Eq. 3.15 into Eq. 3.16. Therefore, projecting data onto the first p principal components will maximise the variance that can be captured in the rank p approximation that was given. This cements the approximating qualities of principal component analysis.

3.3.4 t-Distributed Stochastic Neighbourhood Embedding

In the previous subsection, principal component analysis was introduced as a linear dimensionality reduction method with focus on the capture of global relations to minimise variance. Different non-linear alternatives to principal component analysis exist, collected in the subfield of non-linear dimensionality reduction, or manifold learning [181, 182]. In this subsection the non-linear method of t-SNE (t-Distributed Stochastic Neighbourhood Embedding) is treated.

The t-SNE algorithm was introduced in [183] as a variation of the stochastic neighbourhood embedding algorithm [184], seeking to address the crowding problem [183]. The crowding problem can be coarsely summarised by the observation that the volume enclosed by an n-sphere S^n of radius r scales with the dimension as r^n, and therefore under a dimensionality reduction, the data might become crowded in the lower dimensional representation. As the latent space is typically taken to be $\mathbb{R}^2$ or $\mathbb{R}^3$ for ease of visualisation, this is particularly problematic if the dimension n of the data is large.

The input for the t-SNE algorithm is a similarity matrix of the Euclidean distances between the datapoints, where the data is denoted by $\{x_i\}_{1\leq i\leq m}$. These distances are then converted into conditional probabilities $\mathbb{P}(x_j|x_i)$ as an alternative representation of similarity, representing the likeliness that x_j is a neighbour of x_i. Setting $p_{j|i} := \mathbb{P}(x_j|x_i)$ for short, it is assumed that the $p_{j|i}$ are Gaussian and centered at x_i. This means that to fully specify the probabilities, a standard deviation σ_i will need to be chosen, preferrably in such a way that the variances correlate with local densities of the data. In a densely populated region for instance, the variance should be smaller, as spurious relations might be introduced otherwise. In order to find such values, the perplexity hyperparameter Σ should be set prior to execution of the algorithm, giving an estimate of the expected number of neighbours per datapoint. For

$$H_i = \sum_{j=1}^{m} p_{j|i} \cdot \log_2 \left(p_{j|i}\right) \tag{3.17}$$

the Shannon entropy from information theory measured in bits [185], the value of σ_i so that $\Sigma = 2^{H_i}$ is found by a binary search [136]. Finally, in order to avoid issues with normalisation, $p_{i|i}$ is set identical to 0.

Next, it is assumed that similar distributions $p'_{j|i} := \mathbb{P}'\left[f(x_j)|f(x_i)\right]$ exist for the datapoints under a map f into a lower-dimensional latent space. Here however, it is assumed that the $p'_{j|i}$ are Student t-distributions with one degree of freedom. This assumption is reasonable as these distributions can be realised as infinite Gaussian mixtures, and they are moreover faster to evaluate. Their use also addresses the crowding problem inherited from SNE, that uses Gaussian distributions with shared standard deviation $2^{-1/2}$, as the Student t-distribution is heavier-tailed and therefore allows a wider spreading of the data in the latent space.

If the map f accurately preserves relations, the distributions p and p' should be somewhat similar. This similarity can be measured through the Kullback-Leibler divergence or relative entropy [185]:

$$\mathrm{D}_{\mathrm{KL}}(p_{j|i} \,||\, p'_{j|i}) = \sum_{j=1}^{m} p_{j|i} \cdot \log_2 \left(\frac{p_{j|i}}{p'_{j|i}}\right). \tag{3.18}$$

A cost function can then be defined by summing the divergences over i, and this cost function can be optimised in $\{f(x_i)\}_{1 \leq i \leq m}$ through gradient descent. The output of t-SNE is therefore this sequence of datapoints mapped into the latent space.

3.3.5 Logistic Regression

In the case of linear regression [157], one models a quantitative response by regressing onto a set of predictor variables. This is however not suited to qualitative responses. Instead, in the case of two categories, one can use binary logistic regression to model the probability of the response being of one of the two categories. This method depends on the logit

$$\ln\left(\frac{\mathbb{P}(X)}{1-\mathbb{P}(X)}\right) = \beta_0 + \sum_{i=1}^{n} \beta_i X_i, \tag{3.19}$$

for $\beta_i \in \mathbb{R}$, where the predictors are given by $X = \{X_i\}_{1 \leq i \leq n}$ and $\mathbb{P}$ is a Bernoulli probability measure. The natural logarithm in this equation is a result from exponentiation in the logistic function for the purpose of normalisation, to which the method thanks its name. The logistic function is in fact a sigmoid function, as was mentioned previously in the discussion of neural networks. The model given by the logit can be fit to the data by computing the maximum likelihood estimators $\hat{\beta}_i$.

The exponentiations mentioned above changes the interpretation of the regression coefficients. In linear regression, the coefficient β_i gives the average additive change in the response associated with a unit increase in X_i. In logistic regression, a similar change multiplies the response by e^{β_i}.

It is important to note that when employing regression, multicollinearity can occur. This is the case when several predictors have particularly high correlation, and this can reduce the statistical significance of the regression outcomes [156].

3.3.6 Bivariate Correlation

Correlation coefficients give statistical measures of the relationship between two variables [186, 187]. Although the famous adage says correlation does not imply causation, correlation can still be an extremely helpful tool when exposed to the right level of scrutiny.

For two random variables X and Y with means μ_X, μ_Y and standard deviations σ_X, σ_Y respectively, the Pearson correlation coefficient $\rho_{X,Y}$ is defined as their covariance [186]:

$$\rho_{X,Y} = \frac{\mathrm{Cov}[X,Y]}{\sigma_X \sigma_Y} = \frac{\mathbb{E}[(X-\mu_X)(Y-\mu_Y)]}{\sigma_X \sigma_Y}, \tag{3.20}$$

divided by the product of their standard deviations. For given realisations or datapoints $\{x_i\}_{1 \leq i \leq m}$ and $\{y_i\}_{1 \leq i \leq m}$, statistics will need to be inserted to compute the sample Pearson coefficient:

$$r_{x,y} = \frac{\sum_{i=1}^m (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^m (x_i - \bar{x})^2}\sqrt{\sum_{i=1}^m (y_i - \bar{y})^2}} \tag{3.21}$$

where $\bar{x}$ and $\bar{y}$ represent the sample means. Any correlation measured by the Pearson coefficient is necessarily linear, and due to the normalisation, $r_{x,y}$ will take values in $[-1, 1]$. A value of 0 will mean there is no correlation, whereas an absolute value close to 1 will signify strong correlation.

As an alternative to the Pearson coefficient, the measurements can first be converted to ranked ordinal variables before computing the Pearson coefficient under this transformation. This leads to the Spearman coefficient [186], or rank correlation coefficient:

$$\rho_{R(X),R(Y)} = \frac{\mathrm{Cov}[R(X), R(Y)]}{\sigma_{R(X)}\sigma_{R(Y)}} = \frac{\mathbb{E}[(R(X) - \mu_{R(X)})(R(Y) - \mu_{R(Y)})]}{\sigma_{R(X)}\sigma_{R(Y)}}, \quad (3.22)$$

under the ranking function R, which in turn for a sample as above can be rewritten as:

$$r_{R(X),R(Y)} = 1 - \frac{6\sum_{i=1}^{m}[R(x_i) - R(y_i)]^2}{m(m^2 - 1)}. \quad (3.23)$$

The advantage of the transformation is that the Spearman coefficient can measure monotonic or order-preserving relations, of which linear relations are a particular example. Following this the Spearman coefficient is also less sensitivity to outliers that could skew the linear interpolation in the Pearson coefficient. However, the Pearson coefficient is more powerful, as it incorporates further information such as the sample means, whereas the Spearman coefficient relies only on the rankings.

Chapter 4

Cosmic String Searches

Gravitational-wave searches for cosmic strings are currently hindered by the presence of detector glitches, some classes of which strongly resemble cosmic string signals. This confusion greatly reduces the efficiency of searches. A deep-learning model is proposed for the task of distinguishing between gravitational-wave signals from cosmic string cusps and simulated blip glitches in design sensitivity data from the future Einstein Telescope. The model is an ensemble consisting of three convolutional neural networks, achieving an accuracy of 79%, a true positive rate of 76%, and a false positive rate of 18%. This marks the first time convolutional neural networks have been trained on a realistic population of Einstein Telescope glitches. On a dataset consisting of signals and glitches, the model is shown to outperform matched filtering, specifically being better at rejecting glitches. The behaviour of the model is interpreted through the application of several methods, including a novel technique called waveform surgery, used to quantify the importance of waveform sections to a classification model. In addition, a method to visualise convolutional neural network activations for one-dimensional time series is proposed and used. These analyses help further the understanding of the morphological differences between cosmic string cusp signals and blip glitches. Because of its classification speed in the order of magnitude of milliseconds, the deep-learning model is suitable for future use as part of a real-time detection pipeline. The deep-learning model is transverse and can therefore potentially be applied to other transient searches.

4.1 Introduction

The future Einstein Telescope will have a greatly increased sensitivity compared to the current generation and is expected to detect many more signals, possibly from new sources. Gravitational waves observed thus far have been the product of compact binary coalescences, which are pairs of coalescing stellar- or intermediate-mass black holes and neutron stars [42, 142–144]. Searches, however, are not limited to such systems. One class of unary sources is that of cosmic strings.

Cosmic strings are objects that are conjectured by several theories to have formed in the early universe, and if present, have evolved as the universe expanded [22, 101]. Theoretically they are found through changes in the fundamental group

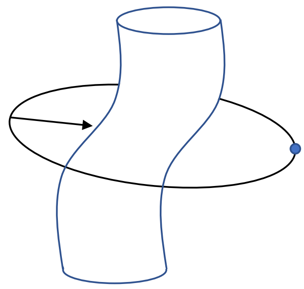

Figure 4.1: A contraction of the circle S^1 to the basepoint on the right, caught on a stringlike deficiency. This deficiency, or defect, does not allow the contraction to complete.

of a vacuum manifold, where a class appears if a contraction of an S^1-embedding fails. This contraction failing due to the presence of a stringlike defect is shown in Fig. 4.1. They should present themselves as strings at cosmological scales. Cosmic strings interact with gravity through gravitational lensing on background light sources due to their angular deficit [101], but also through gravitational waves. The focus of this chapter is the detection of cusps on cosmic strings [188–190]. Cusps can be understood as points on the cosmic string that instantaneously accelerate to the speed of light, and in doing so generate gravitational waves.

As discussed in Subsec. 2.2.3, spontaneous symmetry breaking [100], and therefore the appearance of cosmic strings, may be caused by phase transitions such as the ones associated with grand unification or lower-energy scales. Cosmic strings are therefore of interest to the scientific community as their study can unveil information about both the early universe and a string-theoretical description of the universe [120].

Current searches for cosmic string signatures, of which cusp signals are an example, rely on matched filtering [41, 128, 191–194]. Although these searches have not resulted in observational evidence for the existence of cosmic strings, their results have been used to constrain the model parameters of cosmic strings [128, 192, 194]. Matched filter searches for cosmic strings are hindered by the presence of detector glitches [41, 194], bursts of non-Gaussian noise that may look very similar to modelled cusp signals. Although it is uncertain what glitches will look like in the Einstein Telescope, short-duration glitches that mimic cusp signals are likely to appear. In this chapter, machine learning is employed to demonstrate it is possible to differentiate cosmic string cusps from a common class of transient glitches known as blip glitches in LIGO and Virgo data [195], assuming similar glitches in Einstein Telescope data.

In order to detect cosmic strings, observational signatures are needed. Cosmic strings are massive dynamic objects, producing gravitational waves through a variety of mechanisms. Examples are the formation of cusps and kinks [128].

This work focuses on cusps in closed cosmic strings. The waveform h of such a signal in the Nambu-Goto limit for loop length l at redshift z and tension $G\mu$, in natural units where the speed of light c is taken to be unity, has been computed as a function of frequency f as [128]:

$$\begin{aligned} h_{l,z,G\mu}(f) &= \left[(2/3)^{2/3} 8/\Gamma^2(1/3) \frac{l^{2/3} G\mu}{(1+z)^{1/3} r(z)}\right] f^{-4/3} \\ &\approx \left[0.85 \frac{l^{2/3} G\mu}{(1+z)^{1/3} r(z)}\right] f^{-4/3}. \end{aligned} \tag{4.1}$$

In this formula $r(z)$ is the comoving distance to the loop, or the distance of the observer to the loop, and Γ is the Gamma function. The extrinsic parameters for detection are distance and the sky location. The shape of this waveform in the time domain, and a spectrogram of a strain of noise including this waveform, are shown in Fig. 4.2.

Blip glitches are defined as transient bursts with a duration of around 25 ms with frequency concentrated between 30 and 250 Hz [195]. Depending on the viewing angle and assumptions on loop length [194], a cosmic string cusp signal may occupy this same frequency range. This chapter is focused on the development of methods with respect to blip glitches. However, the methods treated could be extended to any class of short-duration glitches affecting cosmic string cusp searches. Although the morphology of such glitches can differ strongly from cusp signals, in the worst-case scenario they may look near-identical, especially when accounting for the diffusion caused by background noise. This worst-case likeness is demonstrated in Figs. 4.2 and 4.3.

Matched filtering has two major drawbacks. The first is the need for a template bank that sufficiently covers the parameter space which in general can be of high dimension, showcasing issues with scalability. The second is that, specifically for cosmic string searches, matched filtering is not robust to glitches, confusing the two classes due to their similar morphology. These points argue the case that it is worthwhile to explore alternatives to matched filtering for candidate detection in search pipelines. One natural choice is that of neural networks which in theory can address both drawbacks. From a theoretical point of view, it is interesting to note that work is being done towards the replication of matched filtering as neural networks [196]. One could then make a case that neural networks can strictly improve on matched filtering.

This chapter details the training of convolutional neural networks for the task of distinguishing modelled cosmic string cusp signals from artificial blip glitches in simulated Einstein Telescope data. The goal is to both prepare for the arrival of the third generation of detectors, as well as to utilise the higher sensitivity of these detectors to learn about the morphological differences between the two types when obfuscated by detector noise. Having this information may aid in current

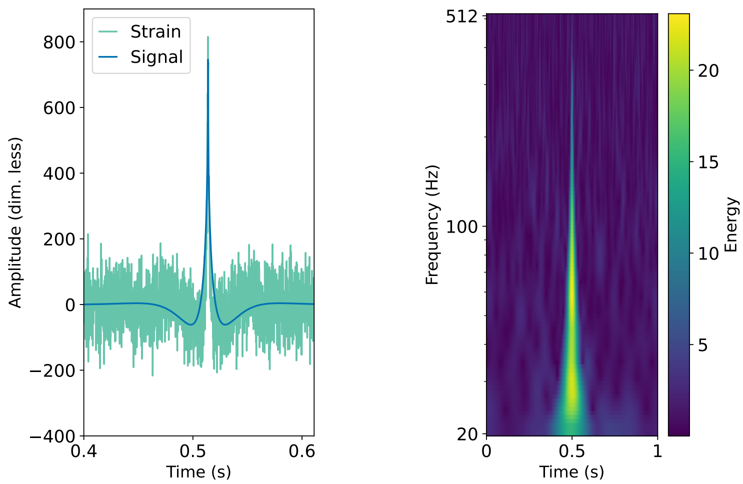

Figure 4.2: A cusp signal with an amplitude of approximately 9.85×10^{-22} Hz$^{1/3}$ prior to injection, overlaid onto the noise it was injected into on the left, and the spectrogram of this strain on the right. Note that the amplitude absorbs the parameters l, z and $G\mu$ per Eq. 4.2.

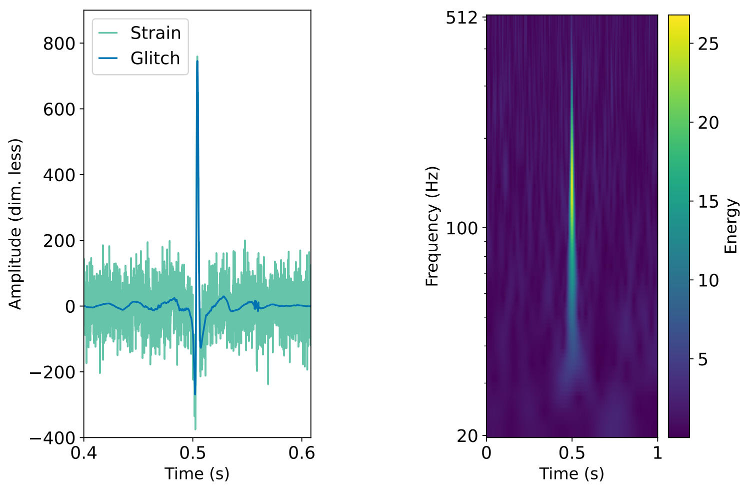

Figure 4.3: A glitch generated by `gengli`, overlaid onto the noise it was injected into on the left, and the spectrogram of this strain on the right. The glitches were scaled to follow the SNR distribution of the cusp signals. This procedure is explained in Sec. 4.2.1.

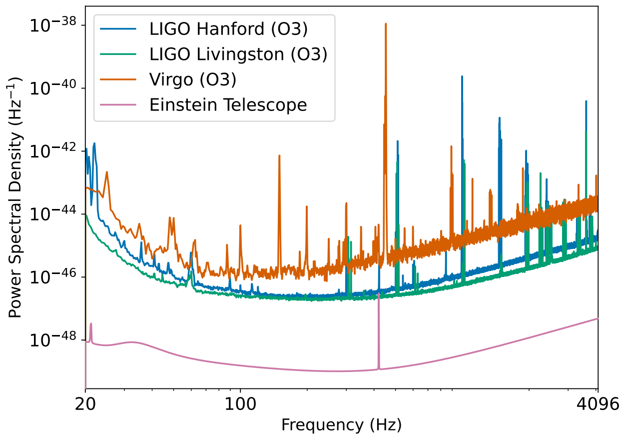

Figure 4.4: The Einstein Telescope design sensitivity compared to the realised sensitivities of the current generation of detectors in the third observational run (O3).

searches in second-generation data as well, as it can be incorporated to design better searches and confirmation tests for observed gravitational-wave candidates.

4.2 Methodology

For the task of training convolutional neural networks on both the as-of-yet undetected cosmic string cusp signals and the per definition unpredictable glitches, a dataset incorporating advanced domain knowledge needs to be constructed. Once this data format is established, the network architecture is treated, along with the design decisions involved. Finally, the methodologies for a comparison to the state-of-the-art and making interpretations of the deep-learning model are described.

4.2.1 Construction of the Dataset

The Einstein Telescope will consist of three detectors in a triangular configuration [197]. As such, three detector strain data streams will simultaneously be collected. In this chapter, these streams will be labeled stream 0 through 2. Einstein Telescope data was simulated by first producing coloured Gaussian noise and then injecting cusp signals and blip glitches into the streams.

For each of the three streams of the Einstein Telescope, a Gaussian noise time series of length twelve seconds was generated, that was subsequently coloured by the PSD representing the Einstein Telescope design sensitivity [198] using `PyCBC` [147]. The design sensitivity of the Einstein Telescope along with the sensitivities of current (second) generation detectors [199] are shown in Fig. 4.4. The noise realisations were then injected with cusp signals to form the positive class, and artificial

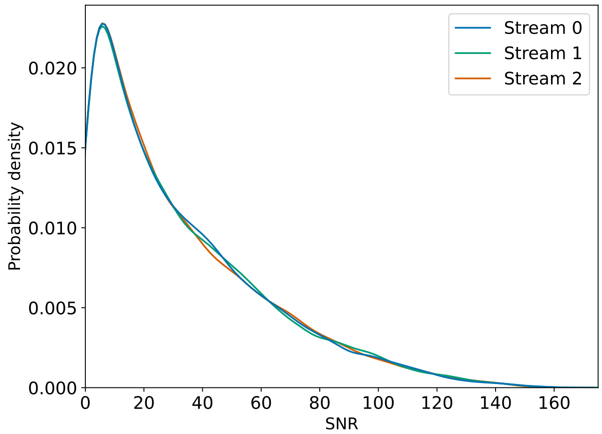

Figure 4.5: The SNR distribution of injected cusp waveforms modelled as a probability density. The streams refer to the streams of the detectors making up the Einstein Telescope.

glitches to form the negative class.

The cusp waveforms in the time domain were generated through the use of the `LALSimulation` package [145]. The function for the generation of the plus-polarised cusp strain components requires three inputs: an amplitude A in $\mathrm{Hz}^{1/3}$ (normalising Eq. 4.1), a high-frequency cutoff f_{high} in Hz past which the waveform will drop exponentially, and a sample period Δ_t in Hz. Here, A represents the amplitudal prefactor in Eq. 4.1:

$$A := 0.85 \frac{l^{2/3} G\mu}{(1+z)^{1/3} r(z)}, \tag{4.2}$$

so that the waveform in the time domain is given by the inverse Fourier transform of:

$$h(f) = A f^{-4/3} (f_{\mathrm{high}} - f)^{+} \tag{4.3}$$

where the plus signifies the taking of the positive part, and the output time series is of this transformed function at sample period Δ_t. In order to randomly generate waveforms, the amplitudes A and cutoff frequencies f_{high} were uniformly sampled from $[10^{-23}, 10^{-21}]$ and $\{20, 21, ..., 4000\}$ respectively. These generated cusp signals, assuming an isotropic distribution, were projected according to the Einstein Telescope antenna pattern and injected into Gaussian-coloured noise sampled at 8192 Hz, before the strain was whitened and cropped to a length of eight seconds. The resulting SNR distributions of this positive class are shown in Fig. 4.5.

The glitches were artificially generated using the `gengli` package [200], which has learned to model blip glitches in the time domain by harnessing generative adversarial networks [201]. Currently, `gengli` approximates the real distribution of glitches in O2 data, specifically that of LIGO Hanford and Livingston. This data includes anomalies which are glitches that are placed in the tail-ends of similarity distributions, and it is therefore possible anomalies showing a different morphology than blip glitches are generated. Using the `gengli` similarity metrics, an accepted region is defined that excludes roughly one in ten glitches that are deemed too dissimilar from blip glitches. These outliers are discarded. This procedure is described in [202].

The true morphology and intensity of Einstein Telescope glitches are currently unknown. It is however reasonable to assume that short-duration glitches similar to blips will be present in the recorded data, and as they are in fact a worst-case scenario in terms of similarity to the cusp signals, they form the best possible preparation. In order to further ensure the robustness of the models to be trained on this dataset, the generated glitches are scaled in amplitude to follow the SNR distribution of the injected cusps shown in Fig. 4.5. This ensures that the models do not learn a difference in SNR distribution. The injection procedure itself differs from that of cusp signals, since `gengli` generates whitened glitches. These glitches are summed as a time series to eight seconds of whitened noise at randomly drawn offsets. The offsets per stream are uncorrelated and the glitches are chosen randomly, meaning there is no detector coincidence for the glitches. Examples of both injected glitches and cusp signals are shown in Fig. 4.2 and Fig. 4.3.

For both classes, no further pre-processing has taken place. In order to both preserve the original information and retain computational efficiency, time series are used instead of alternative representations like spectrograms.

The resulting dataset consists of 30,000 examples (or data points), split into training, validation and test sets of sizes 16,000, 4000 and 10,000 respectively. Each subset is balanced, meaning it is made up of equal parts positive examples (signals) and negative examples (glitches).

4.2.2 The WaveNet Architecture

The convolutional neural networks [159] discussed in this section are implemented in `PyTorch` [174] and were run on the LIGO Data Grid. The specific machine used has the following specifications: Intel E5-2670 CPU, Nvidia Tesla V100 16GB GPU, and 128 GB of memory.

`WaveNet` [203] is an expressive convolutional neural network designed for the generation of high-fidelity speech audio. The architecture is capable of handling long-range temporal dependencies at high sampling rates, achieved by creating a large receptive field through the use of dilated convolutions, or dilations. Dilations allow the network this reception by skipping over a preset number of neurons in each layer, dilating the layers. By appending dilated layers, an exponential increase in the receptive field is gained at the cost of a linearly increasing number

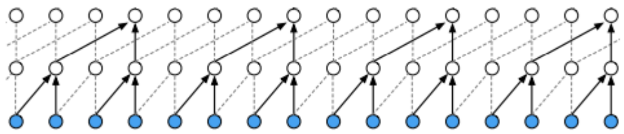

Figure 4.6: Dilation between the layers of the neural network (shown horizontally), retrieved from [203]. As data is passed upwards through the layers, an increasing number of neurons is passed over, creating a larger diagonal reach for the neurons in the top layer.

of layers, as is illustrated in Fig. 4.6.

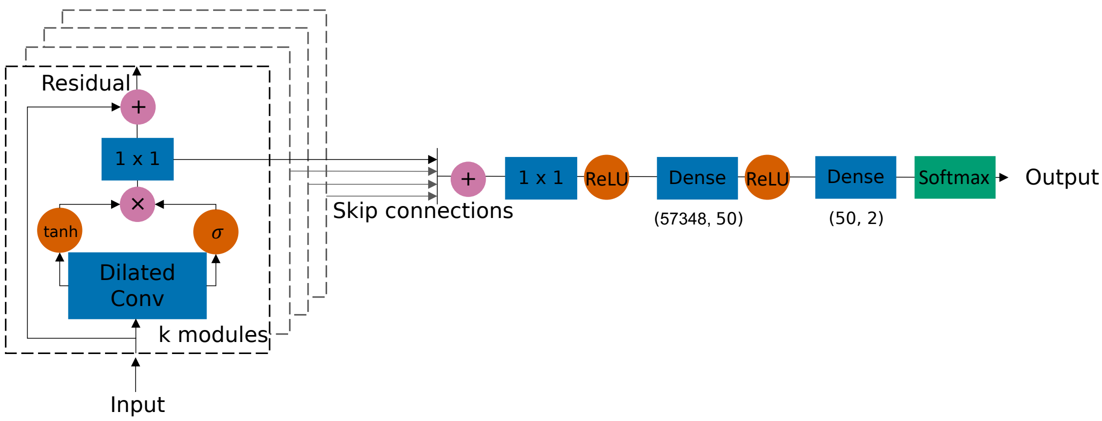

Figure 4.7: Overview of the modified `WaveNet` architecture for a single data stream, adapted from [203]. The hidden layers are coloured blue, the internal activations are shown in orange, and the normalising softmax layer is shown in green.

The major building blocks of `WaveNet` are residual block modules as presented in Fig. 4.7. The figure shows that input to the module is passed through a convolutional layer, after which it is simultaneously passed through both tanh and sigmoid gates. The activations [159] are recombined in elementwise multiplication, where the sigmoid activations modulate the throughput, determining how much of the tanh output activation is passed [204]. The output is convolved with 1×1 filters to reduce the number of parameters before being fed into the residual connection [205, 206]. Note that at this point a copy of the throughput is sent to a skip connection [206].

Inspired by the methodology proposed in [207], where the full `WaveNet` architecture was modified for the discovery of binary black hole systems, modifications have been made for this project as well. The major changes are listed below.

- Instead of encoding the time series amplitudes in a range of 256 possible values (see [203] for details), no such limit is imposed in our implementation;

- The causal structure intended for the dependencies in human speech was removed so as to provide the most possible information to the model;
- The dilated convolutions have a kernel size of 3 to capture fine details, and the dilation within the k-th block module (of 11) is set to 2^k;
- The steps preceding the softmax activations were removed in favour of dense layers. In order to produce a probability, the activations need to be collapsed onto a scalar value in the unit interval. This too is shown in Fig. 4.7.

Together, these changes tailor the architecture to the needs of binary classification instead of the originally intended generation.

4.2.3 Design and Parameter Choices

The first major design choice is the use of an ensemble. Instead of training a single network on the three streams, one network was trained for each, and the three final networks were combined into an *ensemble*. This has several advantages. The first is the handling of different glitches being injected into the streams, therefore not allowing the ensemble to resort to using coincidence for its classification and forcing it to consider morphologies. Second, the independent networks can learn different characteristics during their training phase, averaging out to a more well-informed final decision by the ensemble. This average is taken literally, as the probability $\mathbb{P}(x)$ output by the ensemble for an example $x = (x_0, x_1, x_2)$ of strains is the average of the components networks $\mathbb{P}_i$ for $i \in \{0, 1, 2\}$:

$$\mathbb{P}(x) = \frac{\mathbb{P}_0(x_0) + \mathbb{P}_1(x_1) + \mathbb{P}_2(x_2)}{3}. \tag{4.4}$$

When the time does arrive that coincidence is needed to confirm a candidate detection in joint analysis, these probabilities can be transferred to a central machine instead of the data containing the candidate. This greatly reduces latency, as a single probability is less costly to transmit than a time series.

The weights of each network were determined using stochastic gradient descent, specifically using the `AdamW` optimiser [208] with learning rate 10^{-4} and a weight decay of 10^{-3}. These values were further varied, yielding no significant improvement at this small scale. The batch size was set to 13. Due to the complexity of the model, increases in the batch size resulted in a direct gain in performance, and this trend is likely to continue. For this model, the batch size was limited by memory.

In the training phase, each separate network was trained independently for 20 epochs, resulting in the training and validation cross-entropy losses shown in Fig. 4.8. This phase was repeated multiple times to ensure the optimiser did not get stuck in an avoidable local minimum. It can be read from the validation losses that because of the small batch size, overfitting started between the second and

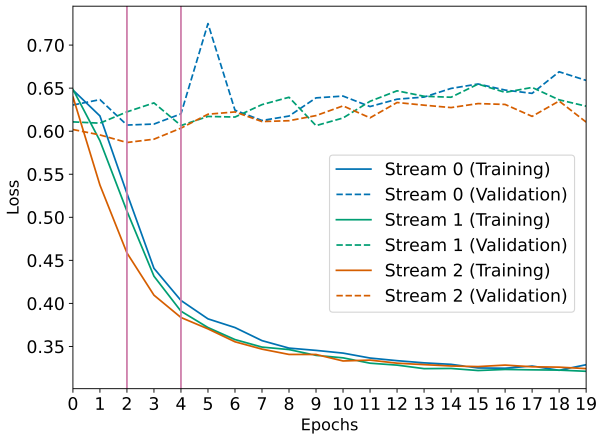

Figure 4.8: Training and validation losses for the three networks. The stopping times are represented by vertical lines, marking the weights used for the networks.

fourth epochs, marked by vertical lines.

From here on, an ensemble is defined by the three ordered epochs at which the training of the networks was halted, denoting the ensemble so created as an $[i, j, k]$-ensemble for i, j, k between the values of 0 and 19. Choosing weights according to the times where overfitting started, the $[2, 4, 2]$-ensemble was established as the initial candidate. Classifiers defined by nearby stopping times in the i, j, k lattice were checked by brute force iteration but gave no improvement over the $[2, 4, 2]$-ensemble. This ensemble was therefore chosen.

Both the individual network thresholds, the ensemble threshold, and combinations of the two were fine-tuned on the validation set. The most important measures used in the fine-tuning were the accuracy, true positive rate and false positive rate. As can be seen from Fig. 4.9, the ensemble probabilities $\mathbb{P}$ on the validation set are highly concentrated in the neighbourhoods of 0 and 1, giving no cause to deviate far from an ensemble threshold of 0.5. The viable range for thresholds to test was set to the uniform set spanning from 0.4 to 0.6 with step size 0.01, with none leading to a significant improvement over the default value of 0.5. A similar line of reasoning has led to thresholds of 0.5 for the component networks.

4.2.4 Comparison to Matched Filtering

A direct comparison between the deep-learning model, a binary classifier, and matched filtering, which is not a binary classifier, is not straightforward. In order to benchmark the deep-learning model against matched filtering, a new balanced dataset of total size 400 was constructed, again following the SNR distribution shown in Fig. 4.5 for both signals and glitches injected into Gaussian noise coloured by the Einstein Telescope design sensitivity. Recall that these were

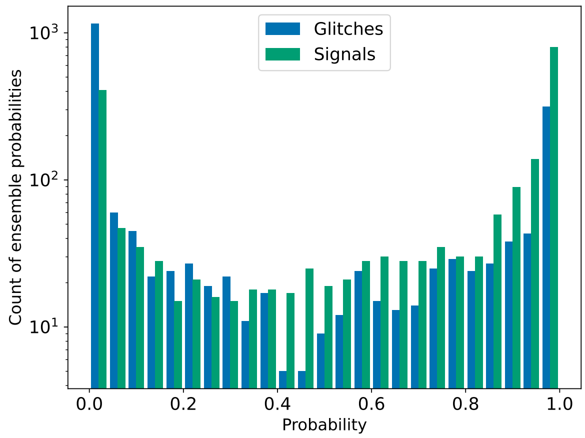

Figure 4.9: Counts of the ensemble probabilities for both classes, with the y-axis in the log scale.

labeled the positive and negative examples respectively.

For a given positive example, each of the three streams was match filtered against the exact injected waveform. This waveform is per definition the optimum filter, and the trigger value is defined as the global maximum of the three SNR time series.

For the negative examples, the optimum filter will not exist as a cusp waveform template, as no true signal was injected. The choice of templates is therefore arbitrary. In order to simulate realistic circumstances, a template bank was created by randomly sampling 20 of the 200 signals generated during the creation of the comparison dataset. Including more templates would be detrimental to the performance of matched filtering, as these additional templates would only allow for the measured SNR to be increased, where it is known no cusp signal is present. Hence, the results from this comparison can be considered conservative. The performance of matched filtering could only be improved by constructing a template bank of cusp waveforms where no template can be matched to blip glitches, which defeats the purpose of the comparison. The remainder of the procedure is identical to that for the positive examples so that the method is internally consistent.

The results of this comparison are presented in Subsec. 4.3.2.

4.2.5 Model Interpretability

Neural networks are notoriously hard to interpret because of their large dimensionality and opaque optimisation procedures. Ideally, however, the discriminative properties the networks have learned would be extracted, in order to better understand the morphological differences between the injected signals and injected

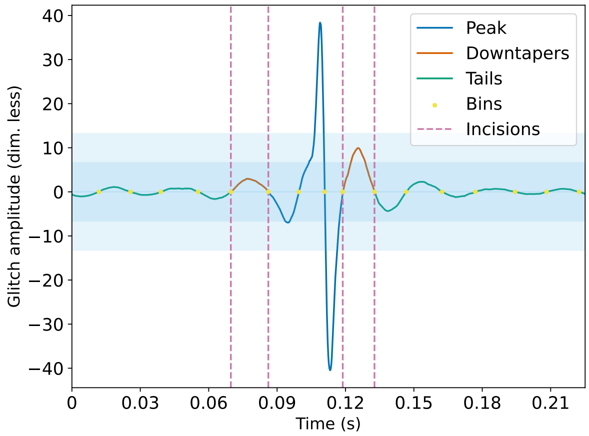

Figure 4.10: The sections of a glitch within a visualisation of the surgery procedure. The shaded areas represent the standard deviation of the amplitude from zero.

glitches. So as to learn what the neural networks have learned, a variety of methods is proposed to interpret the behaviour of our deep-learning model.

Surgeries

The first method employed to better the understanding of our deep-learning model is what will be referred to as *glitch surgeries*. Surgeries are limited to the class of glitches which are not subject to detector antenna patterns through projection, meaning impact can more directly and accurately be measured for glitches. Moreover, they are more readily split into different regions on which surgeries can be performed. The predetermined parts of a selected glitch are excised before reclassifying the modified example and quantifying the change in the ensemble prediction with respect to the original input. In doing so, features salient to the deep-learning model can be identified.

The observation underpinning the procedure is that a glitch $g(t)$ can generally be divided into five regions based on the maximum amplitude within these regions and that these regions together form the sections shown in Fig. 4.10. Note that this subdivision is not unique and other regions can be identified. These regions can be automatically detected by partitioning a glitch into bins delimited by the zero crossings and comparing the absolute maximum of each bin with a function of the standard deviation σ of the glitch amplitude. The edges of the bins are constrained to correspond with zero-crossings to ensure continuity, as an excision amounts to setting the value of the glitch waveform amplitude to zero within the bins that are excised. Whereas continuity is required so the neural networks do not pick up on the transitions, it is not necessary to extend the waveform to be smooth at the bin edges, as this transition is lost within the noise after injection.

First, the *area of peak activity*, which one should note may contain more than one peak, is identified as being the bin containing the absolute maximum amplitude $|\max(g)|$ of the glitch in the time domain. Moving outwards left and right over the bins, starting from the identified bin, the peak area is extended to include adjacent bins if the absolute maximum within these bins exceeds 2σ. The *downtaper* of the glitch starts at the first bin where the absolute maximum within the bin is below 2σ, and the *tail* of the glitch starts at the first bin where the absolute maximum is below σ. Note that these definitions may imply the absence of named sections in a glitch waveform, as for instance, a section corresponding to the downtaper might not exist. This can be the case if the maximum amplitude of the waveform is extremely high compared to the average amplitude. Such glitches can safely be included in the surgery procedure. The excision of a non-existent section amounts to nothing changing at all, and the results from the reclassification will reiterate that the non-existent section did not contribute to the classification.

For a representative sample taken from the dataset, the procedure is then as follows.

1. Choose a glitch example g from the set, and retrieve the ensemble probability $\mathbb{P}(g)$;

2. For any of the three sections, set g identical to zero within the corresponding bins to obtain g_{section} (performing the surgery);

3. Reinject and classify g_{section} before retrieving the ensemble probability $\mathbb{P}(g_{\text{section}})$.

The statistic of interest is then:

$$\Delta_{\text{section}}(g) := \mathbb{P}(g) - \mathbb{P}(g_{\text{section}}). \tag{4.5}$$

Note that this statistic takes values in $[-1, 1]$. The natural interpretation is that a value close to -1 means that the classification has significantly changed, with g being classified as a glitch previously and as a signal following the surgery. A value close to 1 would imply the reverse. A stream from an example with a value of $\Delta_{\text{peak}} \approx -0.47$ is shown in Fig. 4.11. This behaviour can be further explored by considering the changes for the individual component networks within the ensemble.

Activations

Another way of investigating the behaviour of the ensemble is the extraction and visualisation of the *activations* in the hidden layers as a testing example is passed through. This information can then be used to tie certain convolutional filters to specific confusion matrix classes (Fig. 4.12) in the dataset. Note that these are filters according to the terminology of neural networks, not those of matched filtering. Inspiration was drawn from saliency maps [209] from computer vision,

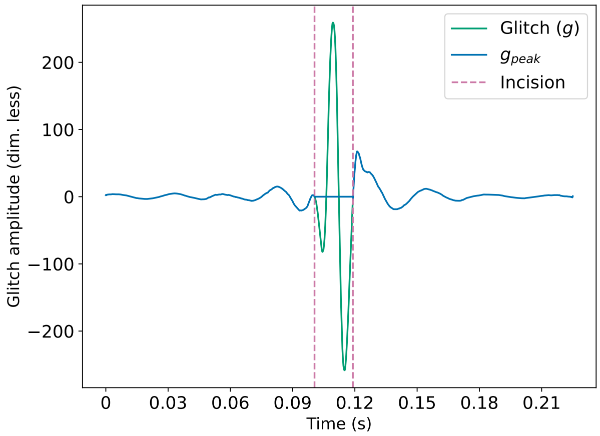

Figure 4.11: A glitch before and after surgery. During the surgery the peak was excised from g, thus obtaining g_{peak}. The corresponding statistic is $\Delta_{\text{peak}}(g) \approx -0.47$.

meant to highlight the most salient and therefore recognisable regions of images. Although projects like `Captum` [210] offer similar ways of interpreting convolutional neural networks, they differ from the method described here, designed specifically for the analysis of time series data.

A straightforward way of obtaining the activations (that also works for general networks) is to deconstruct a given network into an ordered set of individual layers, applying these layers one by one, and saving the outputs before feeding the output forward. Once these values are recovered, the challenge of interpreting the activations is reduced to trying to connect the activation of specific filters to fundamental characteristics of the example that was passed through. This is akin to detailing a collection of neurons that fire when a specific example is seen. As this is an extremely difficult task with high dimensionality, only isolated observations can be made.

In order to understand how the activations can be best visualised, it is useful to review the process of a filter being applied. As a filter is convolved with the one-dimensional time series, a new time series containing a large number of activation values is obtained and fed out. Due to the number of values, a direct plot of the activations would be unreadable. Instead, the values are binned and smoothed with a kernel density estimate. The resulting curve is an indicative visualisation of the activation for the specific filter used. The reader is invited to look ahead at Fig. 4.17, which shows these curves for the examples that will be interpreted in the next section.

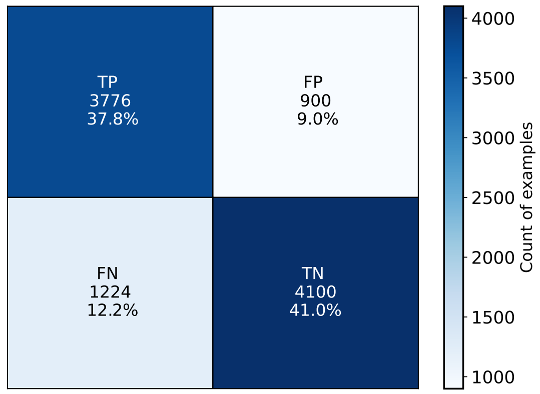

Figure 4.12: The confusion matrix showing the true positives (TP), false positives (FP), false negatives (FN) and true negatives (TN) for the $[2, 4, 2]$-ensemble on the test set, visualised as a heatmap.

Principal Component Analysis

Whereas the extraction of the activations serves mostly to delve into the hidden representations of the data as it is passed through the modules, principal component analysis (or PCA) [211] can be used to analyse the representation in the linear layers. PCA is a dimensionality reduction method that linearly maps vector data into a lower-dimensional space with an ordered basis consisting of what are called the principal components. These components are determined as being the basis vectors carrying the most amount of information measured by variance, and their ordering is based on these same amounts. This means that for instance, the first principal component contributes the most to the overall variance of the dataset. PCA is applied to the second to last dense layer shown in Fig. 4.7, where an input of size $57,348$ is collapsed to an output of size 50 before the latter values are further reduced to a single probability. Based on these numbers one can argue that in this layer the most amount of information is condensed, making it a valuable object of study. In Subsec. 4.3.3 the first two principal components obtained from the dense layer are studied.

4.3 Results

In this section, the numerical results of the chosen deep-learning model are reported and discussed before treating the information extracted from the model by applying the interpretability methods presented in Sec. 4.2.

Metric	**Formula**	**Value**
Accuracy	(TP + TN) / (P + N)	0.7876
True Positive Rate	TP / (TP + FN)	0.7552
False Positive Rate	FP / (FP + TN)	0.1800

Table 4.1: A selection of performance metrics for the $[2, 4, 2]$-ensemble on the test set, along with their formulae. These values were computed from the confusion matrix in Fig. 4.12 using the true positives (TP), false positives (FP), false negatives (FN), true negatives (TN), with P := TP + FP (positives) and N := TN + FN (negatives).

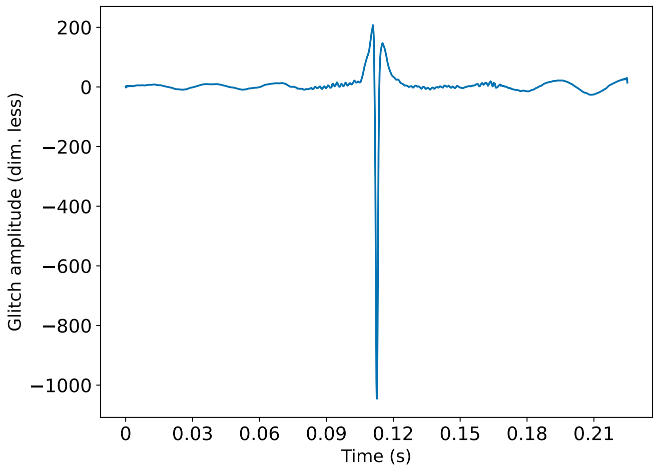

Figure 4.13: An example of a glitch that was falsely classified as a positive with high probability. It is therefore an extreme example of a false positive.

4.3.1 Numerical Results

On the test set, the $[2, 4, 2]$-ensemble yields the confusion matrix visualised in Fig. 4.12, from which the metrics presented in Table 4.1 were computed. In terms of cosmic string searches, the accuracy refers to the model's capability of recovering the injected signals and glitches. Likewise, the true positive rate (TPR) quantifies how well the model can recognise a signal, given that a signal was injected. Finally, the false positive rate (FPR) measures to what degree the model mistakes glitches for signals, given that no signal was injected. This means that a value of 0 is ideal for the FPR, and 1 is ideal for the accuracy and TPR.

Manually investigating the most extreme false positive examples, spurred by the relatively high FPR, most skew high on this metric due to the three streams having been injected with glitches sharing a similar morphology between them, shown in Fig. 4.13. It appears the extremely high maximum amplitude (as compared to the average amplitude) dominates the morphology, leaving the deep-learning model little other distinguishing features to base its classification on. For these

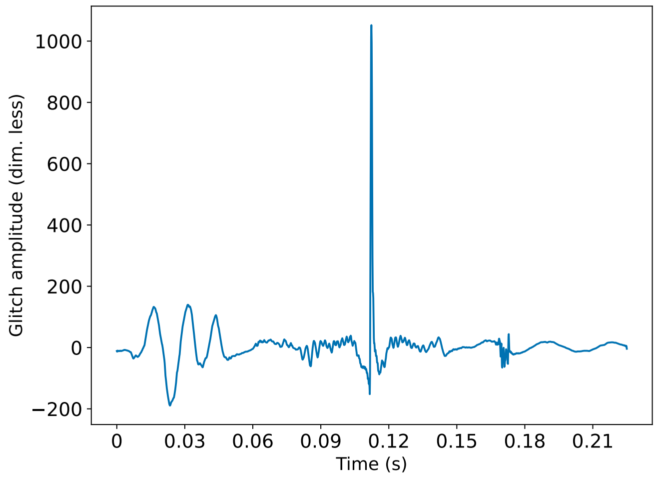

Figure 4.14: A different extreme example of a false positive. One possible explanation for this misclassification is that the component network corresponding to this stream is given more opportunity to identify a signature that is believed to be that of a signal.

specific examples, the model erroneously resorted to a positive classification.

There are a few noteworthy exceptions, each appearing only in one of the three streams that make up an example, showing oscillations in amplitude over a larger period of time. Such a glitch is shown in Fig. 4.14. One possible explanation for these instances is that the component network is given more opportunity to detect the presence of a signal signature and that one such signature is sufficient for the example to be classified as a signal. This underlines the importance of the signal morphology to the deep-learning model.

For the archetypal examples shown in Fig. 4.2 and Fig. 4.3, the component networks for the streams these examples were taken from assigned the glitch a probability of 0.0008 of being a signal, and the signal a probability of 0.7005. This means the networks assign these examples to the right classes with considerable confidence.

Lastly, on the machine used (described in Subsec. 4.2.2), the classification speed of one example (consisting of three data streams of 8 seconds) was computed to be 10 milliseconds on average.

4.3.2 Comparison to Matched Filtering

Conventionally, the performance of matched filtering is measured with positive examples being signals added to noise, and negative examples being drawn from a coloured Gaussian noise background without a signal present. Recall however that the current consideration is the distinguishing power of the deep-learning model and matched filtering for a dataset where the positive examples are injected sig-

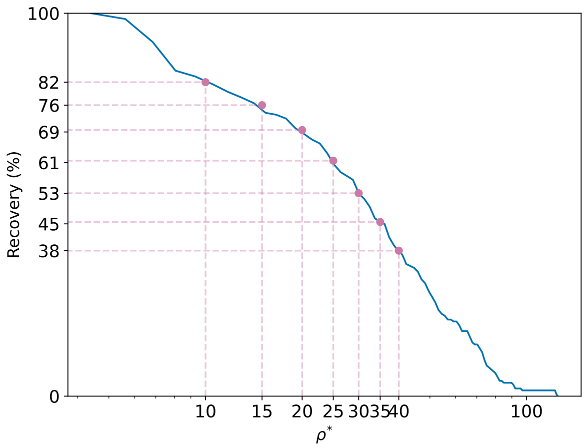

Figure 4.15: The recovery percentage of matched filtering for different values of the SNR threshold ρ^* in log scale. This latter value represents the cutoff from which point onwards an SNR is considered high enough for the corresponding example to be labeled as containing a signal.

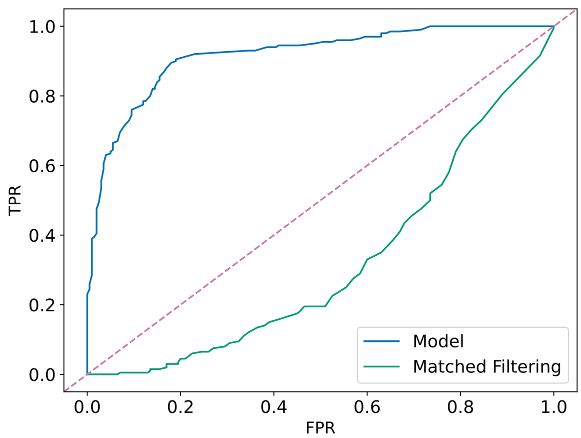

Figure 4.16: The Receiver Operating Characteristic curves for both the deep-learning model and matched filtering on a dataset consisting of injected signals and glitches. The performance of a random binary classifier is represented by the diagonal.

Section	**Average**(Δ_{section})	**Maximum**(Δ_{section})
Peak	0.161	0.832
Downtapers	0.001	0.055
Tails	0.004	0.205

Table 4.2: The average and maximum values of Δ_{section} statistics, separated per section. These values offer a summary of the output from the surgery process, where Δ_{section} measures the impact of a glitch section removal on the model classification.

nals, and the negative examples are injected glitches.

The true positives recovered by matched filtering at different SNR thresholds ρ^* is shown in Fig. 4.15. This choice for the representation of the results was made in order to remain agnostic towards the chosen threshold, which may differ per analysis. A direct comparison between the deep-learning model and matched filtering is shown through the Receiver Operating Characteristic (ROC) curves in Fig. 4.16. The diagonal represents a random binary classifier, meaning that on this dataset, matched filtering is weighed down by false positives to the point where its performance is worse than random classification. In contrast, the deep-learning model is very effective on the same dataset. The conclusion is that the model is better at rejecting glitches than a simple matched-filter search with cosmic string cusp templates by a large margin. A more complete comparison, including for instance the additional mechanisms that would be present in a full gravitational-wave search pipeline and would work to ameliorate false alarms, is deferred to future work.

4.3.3 Interpretability

In this section the results of the various interpretability methods described in Subsec. 4.2.5 for the interpretation of the deep-learning model are presented.

Surgeries

By performing the surgeries described in Sec. 4.2.5, the statistics given in Table 4.2 were obtained. These statistics show that with a division of a glitch into five regions, or three sections per symmetry, the peaks will by far be the most informative, with the downtapers and tails contributing relatively little. The low average and maximum values for the latter section statistics bound the values on the whole set of glitches, indicating that for no negative example the removal of either the downtapers or tails has made a significant impact on the reclassification.

Investigating the outliers near the maximum of 0.205 for Δ_{tails}, it was found that these values stem from glitches with high fluctuations in tails that were removed.

For the values on the low end of Δ_{peak}, a similar observation is made. The removal of the peaks for these glitches left behind fluctuations in amplitude in the downtapers or tails, on which the deep-learning model will presumably base its classification instead. Most of the examples with a high value of Δ_{peak} have very large amplitudes in the peak section, the removal of which confuses the model. An interesting note is that for these examples the network trained on stream 0 seems to be less impacted by the excision of peaks than the other two networks in the ensemble, which is possible evidence that the three networks have learned to identify different glitch signatures. Further manual inspection of the small number of examples with Δ_{peak} near -1, meaning the classification has changed from a glitch to a signal following the surgeries, shows that all have remaining fluctuations in their waveforms. One theory is that the model considers these remnants as the new peak sections, viewing at least one as evidence of a present signal. This observation might suggest that without detecting a clear glitch signature, the model defaults to a signal classification. This would complement the discussion on the false positives in Subsec. 4.3.1.

Relating to the preceding discussion, if the model indeed resorts to analysing amplitude spikes within the sections that remain after a surgery has been performed, this suggests the model considered these sections as secondary to the peak region before. In turn, this suggests that the model does not simply detect rapid changes in amplitude, but has learned to differentiate morphologies.

Activations

Based on the confidence of the deep-learning model, one example was chosen for each of the classes in the confusion matrix, and their activation values were extracted. This means, for instance, that in the case of the true negative, an example with an output probability very close to 0 was selected. The activations of these four examples for a single module are visualised in Fig. 4.17 and each example will be discussed individually.

For the true negative example in Fig. 4.17a, a number of filters show activation, meaning the mean of the density curve is closer to 1 than it is to 0. However, this is with a high spread in the curve, indicating uncertainty in the activation values for this filter. Some filters, such as 27 in green or 37 in blue, show higher certainty. This is however not enough to mislead the model into making a false positive classification.

The filters in the false negative in Fig. 4.17b see barely any activation taking place at all. For this example, there was nothing giving the model the impression there could be a signal present. The only filters showing a semblance of activation do so with little certainty, with output probabilities not high enough to cross the classification thresholds at which point the model would classify the example as positive.

The false positive example shown in Fig. 4.17c seems to invoke response from the

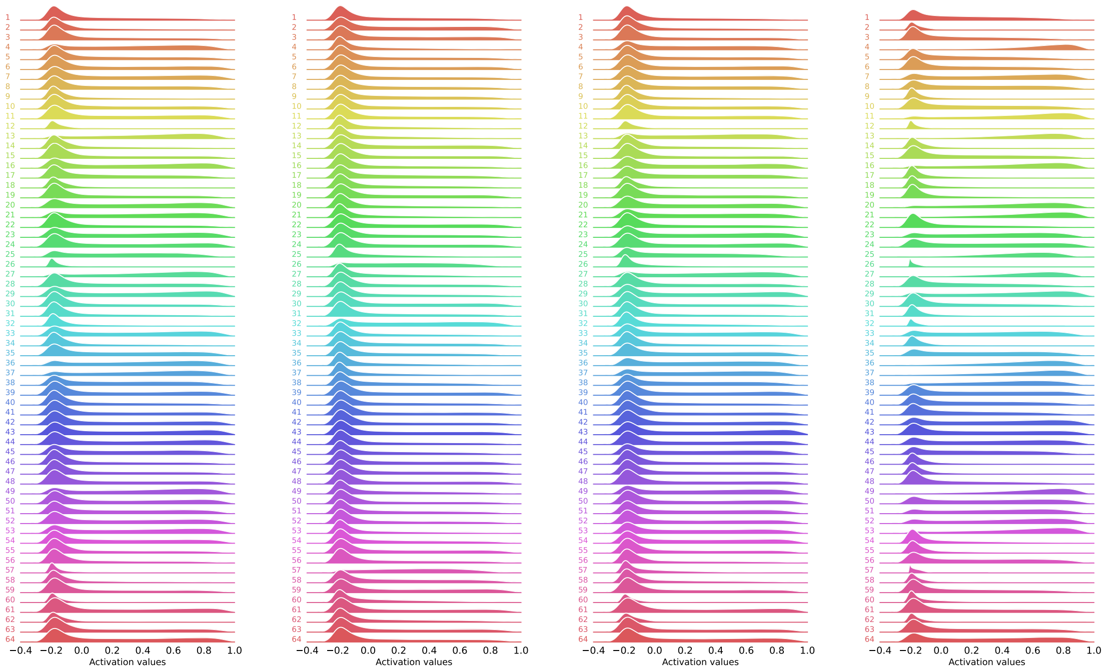

(a) True Negative ($\mathbb{P} = 0.006$)

(b) False Negative ($\mathbb{P} = 0.001$)

(c) False Positive ($\mathbb{P} = 0.982$)

(d) True Positive ($\mathbb{P} = 0.967$)

Figure 4.17: Residual activations output by the 11-th module in Fig. 4.7. Every horizontal axis represents a filter with the y-axis limited to $[0, 0.1]$, showing the distribution of activation values as a density curve. The mode of such a curve being close to 1 indicates high activation values for the filter.

filters, with activity within multiple filters. However, as was the case for the false negative, there is not much certainty. In this case, however, the probabilities did cross the classification thresholds.

Finally, the centroids of the true positive example in Fig. 4.17d skew strongly towards the right, meaning high values of the activations are achieved. A wide array of filters show strong activation, with certainty higher than the previous examples. The individual filters, and the network as a whole, are certain this example contains a signal.

The above observations were made on single examples, and are therefore not guaranteed to generalise. They do however show a clear difference in response to examples from the four classes and are therefore a proof of concept for further investigation. Individual filters can for instance be mapped back to certain sections of the input streams and examined further. This is outside the scope of this work.

As a final remark for this subsection, there are some filters that show little to no activation for any of the four examples, with 9 in orange and $60, 62$ in red being such filters. While this is possibly due to the choice of examples or a lack of need for these filters, it is also possible this is a result of the low number of training epochs, meaning the weights for these filters have not been properly adjusted. If this is the case, one might conclude there is room to improve the model further. One of the ways this could be done is by reducing memory usage during the training phase, leading to better training that may in turn recruit the now dormant filters.

Principal Component Analysis

For the first stream in the test set, the activations of the dense layer were projected onto the subspace spanned by the first two principal components. The results from this projection are shown in Fig. 4.18 and Fig. 4.19, coloured by the probability $\mathbb{P}_0$ output by the first network in the ensemble and the SNR, respectively. These figures are shown in log scale to improve the visual separation between the two classes. Both figures show a portion of the signal population being located in the top left of the plot, whereas a portion of the glitch population is located in the top right. Both classes overlap in the center and are therefore plotted separately to improve visibility. It is relevant to note that in the principal component space, glitches show a larger spread than the signals. This follows from their more varied possible morphologies.

It can be observed from Fig. 4.18 that the probability is related to the first principal component on the x-axis. Compared to Fig. 4.19, the extremes of this same principal component show high values for the SNR. From this, it can be inferred that at least within this representation, the signals and glitches exhibiting the most separability are the ones that are loudest and therefore most obvious to the model. The first principal component can thus be interpreted as a measure of the example class. For the second principal component, there is no such apparent meaning.

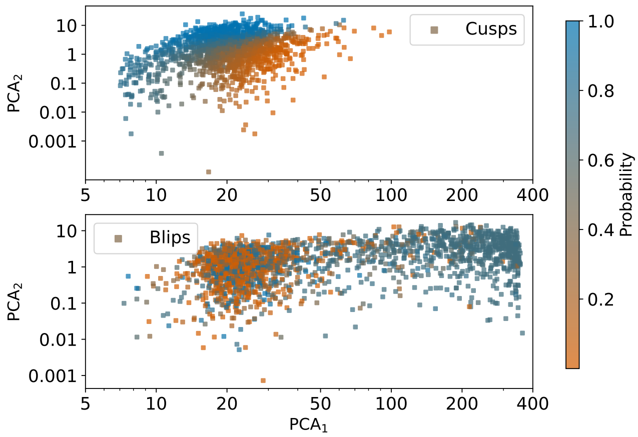

Figure 4.18: A representation of a selection of activations from the dense layer in the space spanned by the first two principal components PCA_1 and PCA_2, in log scale. The plots of the two classes were split to improve visibility that may otherwise be hindered by the overlap of the classes. The points are coloured by the probability $\mathbb{P}_0$ output by the first network in the ensemble.

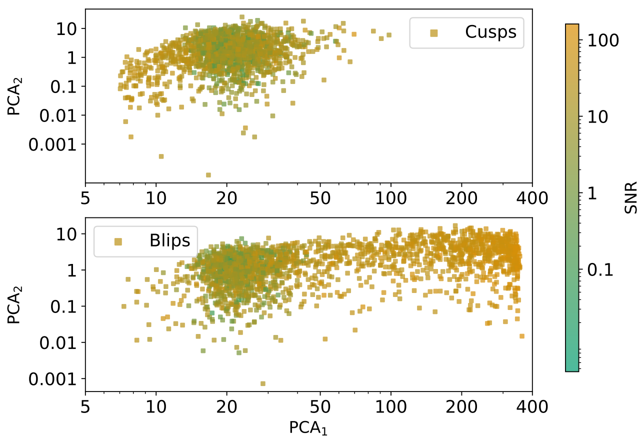

Figure 4.19: A representation of a selection of activations from the dense layer in the space spanned by the first two principal components PCA_1 and PCA_2, in log scale. The plots of the two classes were split to improve visibility that may otherwise be hindered by the overlap of the classes. The points are coloured by SNR.

Chapter 5

Intermediate-Mass Black Hole Binary Searches

Gravitational-wave searches for signals of intermediate-mass black hole binaries are hindered by detector glitches, as the increased masses from stellar-mass systems hinder current generation detectors from observing the inspiral phase of the binary evolution. This causes the waveforms to strongly resemble glitches, which are of similar duration within a similar frequency band. Additionally, precession of the orbital plane of a binary black hole may further warp signal waveforms. In this chapter three neural network-based classifiers for the task of distinguishing between signals and glitches are introduced, with each following different training regimes to study the impact of precession on the classifiers. Although all classifiers show highly accurate performance, the classifier found to perform best was trained following the principle of curriculum learning, where new examples are introduced only after the mastery of easier preceding examples. This classifier obtains an accuracy of approximately 95% on a synthetic test set consisting of signals and glitches injected into coloured noise from the O3 LIGO Hanford power spectral density. The model is compared to matched filtering, the state-of-the-art in modelled gravitational-wave searches, and analysed in search of particular sensitivities to black hole binary parameters. It was found that while the classifier is affected by the total mass of a system, the prediction of a misclassification is most strongly determined by visibility through the signal-to-noise ratio. The analysis of the three classifiers demonstrates that precession is handled differently depending on the training regime, meaning the architecture is not fully robust to precession and advancements can be made through the development of training routines.

5.1 Introduction

Black holes, since the prediction of their existence as solutions of the Einstein field equations in general relativity [31], have been indirectly observed in astronomy through electromagnetic radiation or gravitational interaction with nearby objects [62, 212–216]. The first of such observations came when Cygnus X-1 was discovered as an X-ray source in 1964, which was later recognised as being a black hole [214, 217]. In 2015, the first black hole binary was observed through the messenger of gravitational waves in the event GW150914 [17]. In the years following this first observation, the LIGO-Virgo-KAGRA collaboration has confirmed over

90 detections of gravitational waves over three observational runs [42, 142–144], with the fourth observational run on-going at the time of writing.

Among these detections was the event GW190521 [218]. Before 2020, a mass gap spanning the range between stellar-mass black holes and supermassive black holes [219, 220] was thought to exist, called the range of intermediate-mass black holes. The detection of GW190521 in 2020, however, has placed a black hole resulting from a merger within the intermediate-mass range. Observing black holes within this mass range is of interest as they may serve as precursors to supermassive black holes [221–223], the analysis of which in turn may further the scientific understanding of galaxy formation and evolution [219]. Furthermore, intermediate-mass black holes are considered as possible ultraluminous X-ray sources [224, 225].

The signals of intermediate-mass black hole binaries as measured by current-generation gravitational-wave detectors will be of short duration, as most of the signal falls outside of the sensitive bands. Because of this, the signals can easily be confused with glitches, bursts of non-Gaussian noise that are not of astrophysical origin [195]. Current searches are therefore antagonised by glitches and rely on several types of statistical tests to mitigate them [195, 226, 227]. Despite the systems in place the classification problem remains difficult.

In this chapter the problem of discriminating between gravitational-wave signals from intermediate-mass black hole binaries and glitches is approached using machine learning. The resulting models can potentially be used as additional statistics in searches. In particular, our work uses convolutional neural networks [159]. Convolutional neural networks have previously succesfully been applied to the study of black hole binaries, and more generally compact binaries, as well as to glitches and their separation from such signals [25, 201, 202, 207, 228–240]. In this chapter three such models obtained with different training regimes are introduced, and shown to perform with high accuracies. Other studies typically assume non-eccentricity and no precession of the orbital plane to be present. Precession can distort signals through amplitude and phase modulations [71, 241], which may cause even higher similarity to glitches. In this chapter the robustness of the neural networks to the added effects of precession is tested for the first time. A comparison in performance is made between the neural network models and matched filtering [134], used in current state-of-the-art detection pipelines. This comparison establishes a baseline for the possible improvements enabled by the use of machine learning for this task. The behaviour of the models is interpreted through several data analysis methods in an attempt to reveal differences between the classes of signals and glitches that other methods are not able to probe. Within the class of signals, the impact of precession on model sensitivity is tested and compared to other system variables to both interpret the models and to test their robustness to precession.

Compact binary coalescence describes the coalescence of two compact objects such as black holes or neutron stars in a binary system [34]. In this work, the study of such processes is limited to binary black holes, as neutron stars add the consideration of tidal deformation, and no neutron stars exist within the intermediate-mass range. Stellar-mass black holes are abundantly represented in the current gravitational-wave transient catalogues [42, 142–144], and supermassive black holes too are well known, with the supermassive black hole at the center of the Messier 87 galaxy and Sagittarius A* at the center of the Milky Way having been observed by the Event Horizon Telescope [19, 242, 243]. These observations bound the mass gap of intermediate-mass black holes between 10^2 and 10^5 $M_\odot$ [220]. Proof of black holes with masses within this gap was found only recently with the detection of GW190521 [218], which also holds the record for largest recorded progenitor masses in a binary system detected through gravitational waves. During this event two black holes with masses of 85 $M_\odot$ and 66 $M_\odot$ merged into one black hole with a mass of 142 $M_\odot$, meaning 9 solar masses of the total mass were radiated away in the form of gravitational-wave emission that was subsequently registered by the LIGO and Virgo detectors. Besides establishing their existence, the observation of this event also proves that mergers are a possible formation mechanism for intermediate-mass black holes. It was previously conjectured that black holes in this range could form through mergers, as the formation through gravitational collapse is unlikely [244–248]. An alternative explanation is that these black holes have formed during cosmological inflation, making them primordial black holes [249, 250].

Intermediate-mass black holes are interesting as they may serve as precursors to supermassive black holes [221–223]. A stellar-mass black hole would have to accrete mass at the Eddington limit for a billion years to reach the status of a supermassive black hole [221]. If instead an intermediate-mass black hole were to be taken for the seed, the required time could greatly decrease. The study of intermediate-mass black holes is therefore strongly coupled to that of both stellar-mass black holes and supermassive black holes.

Inspiraling binary black holes may exhibit precession of the orbital plane [71], occuring if the total spin vector of the masses admits anything other than a z-component in the frame where the orbital angular momentum is aligned with the z-axis. Because of this, binary black holes that do not exhibit precession can be referred to as (anti-)aligned spin systems. A binary black hole system (with or without precession) is then defined by 15 parameters. The two black holes each have a mass, with an additional set of six parameters for the spins. There are then two sky location parameters, the luminosity distance, the phase of coalescence, the inclination, polarisation angle, and finally, the time of arrival [71]. With these parameters given, traditionally, the waveform evolution of the system is modelled best by numerical methods [78, 81]. In many applications however, waveform approximants are used. Approximants are models that approximate numerical waveforms to high accuracy at much lower cost [86–92]. In this chapter the waveform model used is the `IMRPhenomTP` approximant [93], a time-domain model that supports precession. Precession manifests itself in the waveform as modulations in the amplitude and phase [71, 241].

Due to the increased masses when compared to stellar-mass binary black holes, the signals of intermediate-mass black hole binaries sweep through the sensitive band of current generation detectors with a shorter duration, allowing for the measurement of only a limited number of wave cycles that start from a lower frequency [220]. This is the result of the waveform duration being inversely proportional to the chirp mass, which increases with increases of the component masses [34]. As a consequence the signal is dominated by the merger and ringdown phases [220], as the inspiral will fall below the low-frequency cutoff [251]. The short duration causes signals to be confused for glitches [220, 226, 227], with blip glitches being a particular example [195]. Blip glitches are defined as transient bursts of non-Gaussian noise that are not of astrophyisical origin, with a duration close to 25 ms and frequencies concentrated between 30 and 250 Hz [252]. This frequency range largely overlaps with the band for intermediate-mass black hole binaries [252]. As an example, in the search for intermediate-mass black holes in data from the third observational run [220], 200214_224526 was mistakenly identified as a signal candidate by the `coherent WaveBurst` pipeline [154]. In truth this turned out to be a blip glitch. Current systems in place to filter out glitches from searches rely on vetoes, gating, and statistical coherence and coincidence tests between detectors [220, 226, 227]. There is however reason to believe these safeguards are not sufficient [220], and precision is required in the analysis of intermediate-mass black hole signals. Their rarity in observation warrants care so that signals are not mistaken for glitches or vice versa, as glitches may impact intermediate-mass black hole population studies.

Aside from unmodelled searches [218, 220], the current state-of-the-art for modelled intermediate-mass black hole searches relies on matched filtering [134, 220]. Efforts have been made to improve on matched filtering in gravitational-wave searches using convolutional neural networks. The process of matched filtering has been approximated using convolutional neural networks [253–255], and work has been done to apply neural networks to searches for compact binary coalescences [207, 228–236]. In addition, research that includes general transient gravitational-wave signals [25, 238, 256] and glitches [25, 201, 202, 237–240] has been done.

5.2 Methodology

The first objective is the training of three classifiers that are able to distinguish the signals of intermediate-mass black holes from blip glitches when both are injected into detector noise. Of the three classifiers, one was trained exclusively on an aligned-spin dataset, one was trained on a precessing-spin dataset, and the third classifier was trained on both datasets through the principles of *curriculum* *learning* [257] and transfer learning [258]. Curriculum learning assumes that a neural network, like humans in education curricula, will be more receptive to the learning of new concepts if they are gradually introduced, starting from easier examples. In the current context, the aligned-spin dataset is considered the easier of the two, meaning the third classifier was trained on the aligned-spin dataset before being transferred to the precessing-spin dataset to resume training. Having three classifiers enables a survey of the architecture robustness to precession.

The second objective is then to compare and interpret the three classifiers, thus testing the robustness of the architecture to the effects of precession, comparing to the state-of-the-art of matched filtering, and dissecting the behaviour of the models.

From this section onwards, the aligned-spin examples, precessing-spin examples and glitch examples will be referred to as the A, P and G-sets respectively. The resulting classifiers trained on these examples will be labeled $\mathcal{A}$ for the aligned-spin classifier, $\mathcal{P}$ for the precessing-spin classifier, and $\mathcal{C}$ for the curriculum learning classifier. This leads to six different pairings of classifiers and datasets, as the glitches are always assumed to be included, and every classifier will be interpreted as a function from the datasets into the unit interval $[0, 1]$. The pairings are shown ahead in Table 5.1.

In this section this methodology will be covered, split over four subsections. After treating `WaveNet`, the generation of the datasets is treated. Next, the training on these datasets is explained. Lastly, the methods for the evaluation of the results are described.

5.2.1 The WaveNet Architecture

Originally designed for the generation of audio speech waveforms, `WaveNet` is a convolutional neural network [159] architecture designed to learn sequential data characteristics [203]. The architecture achieves this by using causal convolutions, which impose a canonical time direction on the data, and allow for better predictions to be made. Such concepts are widely used in speech recognition, where recurrent neural networks have been shown to be highly effective [259] by directing the flow of information between layers. The `WaveNet` architecture uses dilation in its causal convolutions to capture this effect. Dilations allow the network to have a wide receptive net reaching back through its layers at a limited cost, as information is dilated by removing neurons from the feed-forward loop. The name of the dilations is derived from the increasing number of neurons that are disregarded in every subsequent layer.

The `WaveNet` architecture has previously seen success both in the generation of triggers for intermediate-mass black holes [207] and the classification problem of distinguishing between cosmic string cusp signals and blip glitches [25]. In the current work the architecture from [25] is used, with only the dimensionality of

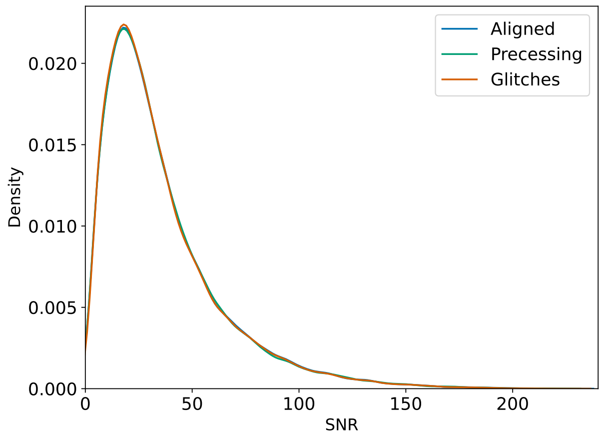

Figure 5.1: The SNR distributions of the aligned spin waveforms, precessing spin waveforms and glitch waveforms. The datasets were generated in such a way that each purposefully follows the same distribution.

the dense layer being changed, as the throughput data will be of a lower sampling rate. The architecture is therefore adapted from waveform generation to binary classification on datasets generated by current-generation gravitational-wave observatories. The neural networks are implemented in `PyTorch` [174] and run on the LIGO Data Grid, specifically on an Nvidia A30 GPU.

5.2.2 Data Generation

Three types of data were generated. The first types are the waveforms of the intermediate-mass black hole binaries, consisting of signals with either aligned spins, which form the first type, or precessing spins that form the second type. The injection tables containing the parameters for the binaries were generated using the `inspinj` function from `lalapps` [145], resulting in $50,000$ parameter samples of both types. The parameter distributions were set to uniform save for the distance, which was set to be logarithmically distributed on $[200, 220]$ Mpc. This range is chosen to target a specific distribution in SNR. The specified total mass range was set to $[150, 600]$ $M_\odot$ with mass ratio within $[1, 10]$ and spin magnitudes within $[0, 0.99]$. The low frequency cut-off was set to 10 Hz. The inclusion or exclusion of precession is regulated through a flag that forces aligned spins when set. The waveforms for the rows in the injection tables were generated in the time domain using the `IMRPhenomTP` approximant [93].

For the third type a total of $50,000$ blip glitches were generated using the `gengli` package [200, 202]. The `gengli` package strives to characterise the O2 blip glitch population of LIGO Hanford and Livingston, allowing the generation of unique new blip glitches through a generative adversarial network.

The two signal types were injected into noise using `PyCBC` [147]. For every signal waveform, four seconds of Gaussian noise was generated at a sampling rate of 4096 Hz. The positive class A of strains including aligned spins was then constructed by projecting the waveforms onto the LIGO Hanford detector and injecting them into the noise at randomly drawn times, and the same was done for the precessing-spin waveforms to form the second positive class of strains P. Colouring was done using the realised O3 Hanford PSD [199], based on the first three months of O3 data [260]. At the end of the injection procedures the strains were whitened.

Glitches were injected following a similar procedure to form the negative class G of glitch strains, with the difference that the glitches output by `gengli` are output in whitened form.

During the injection procedures, the optimum SNR was computed for all examples. The resulting distributions are shown in Fig. 5.1.

Additionally, `gengli` allows for the computation of similarity metrics between two glitches. These metrics are the Wasserstein distance defined as a distance between probability distributions [261], the match, and the cross-covariance [202]. In order to assign every individual glitch a value on these bivariate metrics, 400 typical blip glitches were identified and selected as reference points. In this way the distance metrics can be reduced to a univariate function. For each glitch and each metric, the metric was computed for the glitch and every one of the references, and the average was then taken. The result is one average value per glitch and metric. Small values of these metrics imply expected glitch morphologies, with higher values corresponding to glitches that deviate from the norm.

From these strains two balanced datasets $A \cup G$ and $P \cup G$ with a total size of $100,000$ examples each were created, sharing the same glitches to allow for better comparisons, in particular in the learning phase. These datasets were split into balanced training, validation and test sets of sizes $70,000/10,000/20,000$. Effectively, every classifier is therefore called on three equally sized sets of examples: the aligned-spin examples in A, the precessing-spin examples in P, and the glitch examples in G.

5.2.3 Training Regimes

As in preceding work [25], the networks were trained using stochastic gradient descent using the `AdamW` optimiser [208] with a learning rate of 10^{-4} and weight decay of 10^{-3}. Tuning of these parameters was shown not to improve results. The batch size was set to 30, and each classifier is trained for 20 epochs before interpreting the training and validation cross-entropy losses. As the classifiers are functions from the data into the unit interval due to the final softmax layer, a threshold for the classification needs to be chosen.

In the case of the curriculum classifier, the training phase is started on the aligned-spin set $A \cup G$, before transferring the classifier to continue the training phase on the precessing-spin set $P \cup G$ without overfitting visually setting in. Through this process, it can be measured how much is left to learn from the precessing-spin examples after having learned the aligned-spin examples.

5.2.4 Classifier Evaluation

In order to probe the behaviour of the resulting classifiers, three data analysis methods were employed. These methods will be described in this subsection, and were used to compare the classifiers both to matched filtering and to each other to infer how the different training recipes have affected the classifiers.

The first method is that of principal component analysis (PCA) [157]. PCA is a linear dimensionality reduction method that linearly projects data onto a subspace spanned by an ordered basis of vectors called the principal components, where the ordering in the basis is determined by the explanatory power of the basis elements as measured by variance. As such, lower-dimensional approximations can be constructed by ommitting basis elements at the cost of the variance being reconstructed. In this work, the PCA implementation in `scikit-learn` [262] is used to investigate the sensitivity of the classifier to waveform parameters.

The second method is an alternative dimensionality reduction called t-distributed stochastic neighbourhood embedding, or t-SNE [183]. Whereas PCA is linear, t-SNE relaxes the assumption of linearity, and is therefore part of a class of algorithms called non-linear dimensionality reduction methods, or alternatively, manifold learning [181, 182]. In contrast to PCA that captures global relations to minimise variance, t-SNE attempts to balance local and global relations, guided through a hyperparameter called the perplexity. The perplexity can be interpreted as an estimate of the number of neighbours each datapoint should have. The algorithm then operates by converting the Euclidean distances in the high-dimensional space to probability distributions that represent pairwise similarities, and it is assumed that similar distributions exist in the lower-dimensional space of latent variables. Typically this algorithm is used for visualisation and the latter space is set to be two-dimensional. If these distributions can be similarly realised, then the dimensionality reduction must be possible and successful. Therefore, t-SNE attempts to maximise this similarity by minimising the Kullback-Leibler divergence [263] between the distributions using gradient descent. This algorithm is also implemented in `scikit-learn`, and is used in this chapter to search for clusters of misclassifications in the waveform parameter space.

The third and last method is that of bivariate correlation, in particular through the Pearson and Spearman correlation coefficients [187]. The Pearson coefficient is a normalised measure of the covariance that can detect linear relations. As a covariance-based statistic it is parametric in the statistical sense, as it requires assumptions to be made on the distribution. Alternatively, the Spearman coefficient

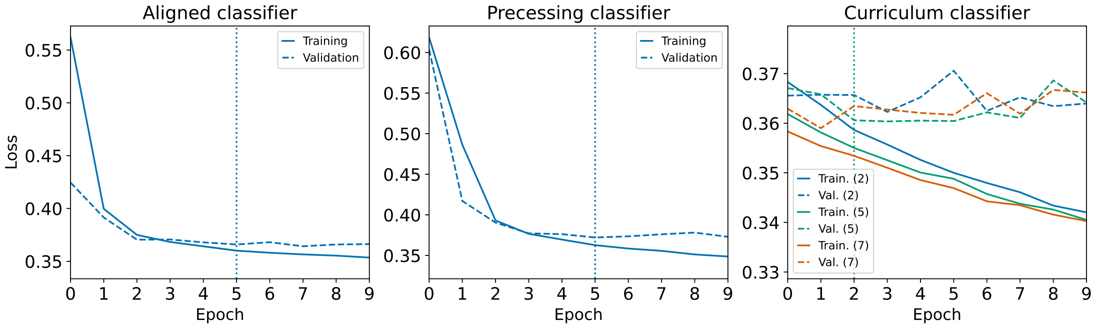

Figure 5.2: The training and validation losses for the three classifiers, with the stopping times shown as a vertical dotted line. For the curriculum classifier, the losses for three different stopping times of the aligned-spin classifier (which were then used as the starting point for the training on the precessing-spin dataset) are shown. Since the classifier ultimately chosen stopped its training on the aligned-spin dataset at the fifth epoch, the stopping time is shown in green accordingly.

is computed in a similar manner, but using ranking statistics. It is therefore non-parametric and moreover only assumes monotonicity, meaning this coefficient can capture non-linear relations. As there is no a priori reason to assume either coefficient will perform better for this work, both are used to investigate the correlation of classifier outputs with different waveform parameters and derived quantities.

5.3 Results

Relying on the methods described in the methodology, this section presents the results, starting with the proceedings from the training phases and model selection. Once the models have been established in the first subsection, a comparison is made to matched filtering in the second subsection, and in the last two subsections the input data is related to the classifier outputs. This is done first by investigating the extreme examples in the false classes of the confusion matrix, and secondly by testing the sensitivity of the classifiers to the different waveform parameters. This relates the misclassifications to the parameters in both directions. By comparing the classifiers during these analyses the robustness against precession is evaluated.

5.3.1 Model Selection

The training and validation losses for the three different classifiers are shown in Fig. 5.2. Even though each classifier was trained for 20 epochs, overfitting had set in before training had reached the last epoch, and therefore only the results of the first 10 epochs are shown for each. Although both the true positive rate and false positive rate are optimised, if in the following discussion two classifiers show equal overall performance, the classifier with the lowest true positive rate will be favoured. The reason for this is that a pipeline potentially depending on this model for triggers would likely rather veto false positives than miss true positives.

In the case of the aligned-spin classifier, the classifiers corresponding to epochs 2, 5 and 7 appear to be the most promising candidates based on the training and validation losses. Comparing statistics computed from the confusion matrix, it turns out that the true positive rate is identical for the three, but the false positive rate is lowest at epoch 5. Although the absolute difference is not much, the accuracy is highest in epochs 5 and 7. Therefore, epoch 5 was chosen as the stopping time.

Following identical reasoning for the precessing-spin classifier, based on the losses, the best stopping times appear to be either epoch 3 or 5. Comparing on the same metrics as before, although the false positive rate is comparable for both classifiers, the true positive rate and accuracy are highest for epoch 5. Therefore here too, epoch 5 was selected as the stopping time.

Three training phases were started for the curriculum classifier, corresponding to the stopping time candidates for the aligned-spin classifier, which were used as the starting times. The reason for starting three such phases is that it is unclear to what degree the knowledge of precession should supersede the knowledge that has already been extracted from the aligned-spin dataset, as the former can serve either to overwrite or as complement. The different phases lead to three candidates, identified by ordered pairs of the stopping time for the aligned-spin classifier (or starting time for the curriculum classifier) and the stopping time of the curriculum classifier: $[2, 3]$, $[5, 2]$ and $[7, 1]$. For all three, the corresponding classifiers obtain approximately equal local optima in their respective accuracies. The true positive rates and false positive rates are correlated and of comparable values on an absolute scale, meaning a variation in epoch that leaves the local optimum in accuracy is not rewarded with a better trade-off between these rates. Based on the classification problem at hand, the classifier with the highest true positive rate is chosen, which is $[5, 2]$.

Tuning of the treshold has lead to no improvements, as the outputs of all classifiers are heavily concentrated on the outer edges of the unit interval. The treshold is therefore set at 0.5, with the output being rounded to a positive classification label of 1 if and only if the output strictly exceeds 0.5. Per complement, the output is rounded to a negative classification label of 0 otherwise.

The confusion matrices for the three classifiers on the test sets are shown in Table 5.1, and the statistics computed from these matrices are are shown in Table 5.2. It can be seen from the latter table that the true positive rate for the sets $A \cup G$ and $P \cup G$ is highest for the curriculum classifier, although the false positive rate is also highest for this classifier. Regardless, the accuracy is highest for the curriculum classifier on both the aligned-spin and precessing-spin datasets. This is also supported by the area under the curve (AUC) values of the receiver operator

Set	TP	FN	TN	FP
$\mathcal{A}[A \cup G]$	9144	856	9851	149
$\mathcal{A}[P \cup G]$	9097	903	9851	149
$\mathcal{P}[A \cup G]$	9037	963	9783	217
$\mathcal{P}[P \cup G]$	8993	1007	9783	217
$\mathcal{C}[A \cup G]$	9332	668	9762	238
$\mathcal{C}[P \cup G]$	9254	746	9763	237

Table 5.1: The confusion matrix of the three classifiers over the two test sets, using the notation introduced in Sec. 5.2. The table shows the true positives (TP), false negatives (FN), true negatives (TN) and false positives (FP).

Set	Accuracy	TPR	FPR
$\mathcal{A}[A \cup G]$	0.9498	0.9144	0.0149
$\mathcal{A}[P \cup G]$	0.9474	0.9097	0.0149
$\mathcal{P}[A \cup G]$	0.9410	0.9037	0.0217
$\mathcal{P}[P \cup G]$	0.9388	0.8993	0.0217
$\mathcal{C}[A \cup G]$	0.9547	0.9332	0.0238
$\mathcal{C}[P \cup G]$	0.9509	0.9254	0.0237

Table 5.2: The metrics of the three classifiers over the two test sets, using the notation introduced in Sec. 5.2. These metrics were computed from the confusion matrices as $(\mathrm{TP} + \mathrm{TN})(\mathrm{TP} + \mathrm{FN} + \mathrm{FP} + \mathrm{TN})^{-1}$ for the accuracy, $\mathrm{TP}(\mathrm{TP} + \mathrm{FN})^{-1}$ for the true positive rate (TPR), and $\mathrm{FP}(\mathrm{FP} + \mathrm{TN})^{-1}$ for the false positive rate (FPR).

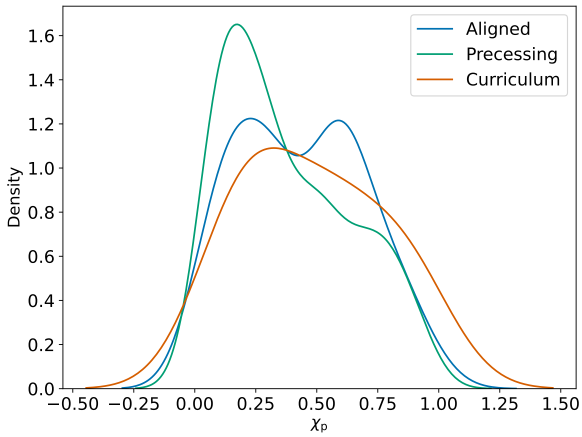

Figure 5.3: The distribution of the false negatives unique to each of the three classifiers over the values of the effective precessing spin parameter χ_p. This plot shows that the false negatives are affected by precession in ways unique to each of the classifiers.

characteristic (ROC) curves, evaluating to 0.9486 for the aligned-spin classifier, 0.9399 for the precessing-spin classifier, and 0.9527 for the curriculum classifier. The curriculum classifier can therefore be appointed as the classifier that performs best, beating the classifiers that were trained on exclusively aligned-spin or precessing-spin examples.

Between the aligned-spin and precessing-spin classifiers, the first achieves a higher accuracy on the set $P \cup G$ than the latter does. This is remarkable. Closer inspection learns that this stems from the FPR. Although the TPR are close together, the FPR for the precessing-spin classifier is higher on $P \cup G$, affecting its accuracy. This means the precessing-spin classifier is more confused by glitches, and that this is the main driver for its weaker performance in comparison.

Returning to the curriculum classifier, it is difficult to discern what it is exactly that the curriculum classifier has learned that makes it outperform the other classifiers on the positive examples, in particular why it performs better on the aligned-spin dataset than the aligned-spin classifier does. However, these classifiers share many false negatives. This observation will be revisited in the later analysis in this section. When taking the false negatives unique to each classifier and considering their distribution over the effective precessing spin parameter χ_p [73] (Subsec. 2.1.6) computed for the injected waveforms, a clear difference can be seen, as shown in Fig. 5.3. For the precessing-spin classifier, the false negatives are concentrated at a lower value of χ_p, indicating that this classifier might have accurately learned the effects of high precession to the waveforms. For the aligned-spin classifier, the distribution overtakes that of the precessing-spin classifier, meaning this classifier has comparatively more difficulty at higher values

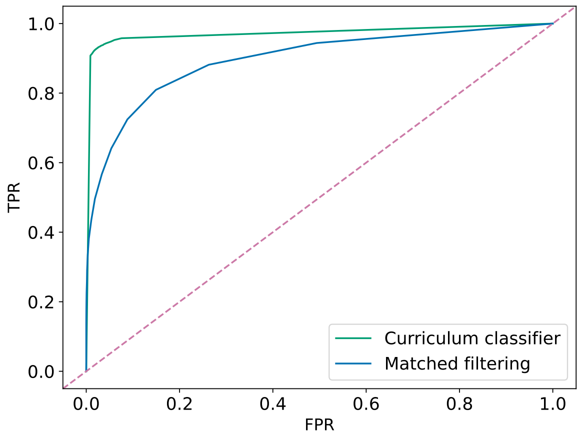

Figure 5.4: The Receiver Operating Characteristic curves of the curriculum classifier and matched filtering evaluated on the held-out test set. The diagonal represents a random classifier.

of χ_p. This is a natural consequence of precessing-spin data having been withheld from the aligned-spin classifier. The distribution for the curriculum classifier looks more evenly distributed, which could suggest that this classifier has succesfully absorbed the information it was provided by the curriculum.

5.3.2 Comparison to Matched Filtering

In order to compare the classifiers to matched filtering, the performance of both methods was evaluated on the test set $A \cup P \cup G$. It should be noted that this comparison is only informative for this comparison, and does not allow conclusions to be drawn for matched filtering based searches in general. In a pipeline there would be mitigation of false alarms, for instance through coincidence in the network of detectors.

For both the aligned-spin and precessing-spin examples, the choice was made to use the injected signal as the template. This means the previously computed optimum SNR can be used, and also means there is no possible disadvantage for matched filtering in this process. For the glitch examples no theoretical optimum filter exists, as these examples do not contain a signal in the first place. Instead, a small bank of 20 templates was constructed, consisting of 10 templates uniformly sampled from the aligned-spin waveforms, and 10 templates uniformly sampled from the precessing-spin waveforms. This bank was then ran against the glitches in the test set. The use of a larger template bank would be unfair to matched filtering, as the template bank puts a lower bound on the SNR that matched filtering will find. The inclusion of other templates could therefore only increase the SNR, leading to more false positives, as any trigger produced is per definition erroneous.

The ROC-curves for both the curriculum classifier and matched filtering obtained on the test set are shown in Fig. 5.4. This figure shows that the performance of matched filtering is strictly dominated by that of the curriculum classifier. In line with the statistics from the previous subsection, this classifier performs extremely well on the test set.

For the aligned-spin and precessing-spin classifiers, the ROC-curves are highly similar, and are therefore not shown. This is supported by the similarity in AUC-values given for the three classifiers in the previous subsection. In terms of robustness, little can be deduced from a comparison to matched filtering. Due to the very similar performance, and the objections to an increase in the size of the template bank used, no statistically sound conclusions can be drawn for the signals.

For the glitches an interesting observation is revealed. Computing the 90-th percentile of the SNR values within the glitches gives an SNR threshold of $\rho^* \approx 16.43$, which can be used to set a bound for the most confident false positives of matched filtering. Comparing the average value of the glitch metrics within these false positives to the average values over all glitches, the cross-covariance and match do not differ significantly. The value of the Wasserstein metric however, which on all glitches ranges between 1 and 120 with an average of 17.66, is on average 48.76 for the false positives of matched filtering with ρ^* as the threshold. As large values of the Wasserstein metric imply anomalous glitches [201], matched filtering seems particularly sensitive to glitches deemed anomalous on this metric. This is supported by the lack in difference of the average match scores, as the match is computed based on the same inner product that is used by matched filtering in the computation of the SNR.

The comparisons in this subsection show that on the test set, matched filtering is strongly outperformed by the three classifiers, and that this can possibly be predicted by the values on the Wasserstein metric.

5.3.3 Extreme Examples

In this subsection extreme examples of misclassifications are considered in order to deepen the understanding of the classifier behaviours through these pathologies. A pathological example is defined for this purpose as either a signal or a glitch example that contains a high SNR injection that should be clear to the classifier but is not, as evidenced by a classifier output far-removed from the true label of the strain. An example would be a signal with high SNR for which the strain with true label 1 is assigned a classifier output near 0. This definition will be used in this subsection. The pathological examples should in theory capture the morphologies that are most confusing to the classifiers. Studying these examples should therefore reveal information on parameters and statistics that lead to extreme misclassifications made by the models, and can be used to direct further investigation of specific parameters, as will be done in the next subsection on parameter sensitivity. It is also a vehicle for discovering differences between the three classifiers, allowing to survey the robustness.

A number of interesting pathological examples surface. The aligned-spin classifier and curriculum classifier share one false negative in the aligned-spin set that has a particularly low classifier output, even when compared to the other extreme false negatives. The precessing-spin classifier, however, assigns this same example an output near 1. When comparing the injection values for this example to the average values in the false negatives, it becomes apparent that the masses and inclinations are much higher at a total mass of approximately 592 $M_{\odot}$, and inclination of 2.99 respectively. This places this particular example at the outer edges of the sample ranges, likely causing the classification differences.

Furthermore, there is one false negative in the aligned-spin set that is shared by all three classifiers. This example has a very high primary mass of 497 $M_{\odot}$ but a relatively low secondary mass of 52 $M_{\odot}$. This leads to a high mass ratio and chirp mass. The signal is therefore of short duration, which is a probable cause for the confusion of the three classifiers.

In general the aligned-spin classifier and curriculum classifier share many examples. Further investigating, most false negatives show a high primary mass, but the secondary mass is mostly higher for the aligned-spin and curriculum classifiers than they are for the precessing-spin classifier. This marks the mass ratio as a property of interest for general misclassifications, as well as for tests of robustness. Cross-comparing, both the aligned-spin classifier and curriculum classifier did fairly well on the false negatives of the precessing-spin classifier. This leads to the observation that the former two classifiers have learned signatures that likely save them from making the same mistakes as the precessing-spin classifier. The curriculum classifier performed well on the false negatives of the aligned-spin classifier as well, suggesting that within the confines of these extreme examples, this classifier has eliminated at least some of the difficulties that plague both the aligned-spin and precessing-spin classifiers.

The inclination is also a noteworthy injection parameter. On the pathological false negatives of the precessing-spin dataset the inclination value is on average lower for the aligned-spin classifier (1.21) and curriculum classifier (0.9) than it is for the precessing-spin classifier (1.55), meaning more of the pathological false negatives of the precessing-spin classifier are closer to being edge-on. This difference is not present to the same extent in the false negatives of the aligned-spin dataset. One possible explanation could be that the inclusion of precession has made the classifier more sensitive to modulation, considering the inclination as a constant modulation.

In the case of glitches, the false positives have a large intersection, with only the precessing spin classifier making unique mistakes. There is however nothing in terms of the glitch similarity metrics that could distinguish these mistakes. Al-

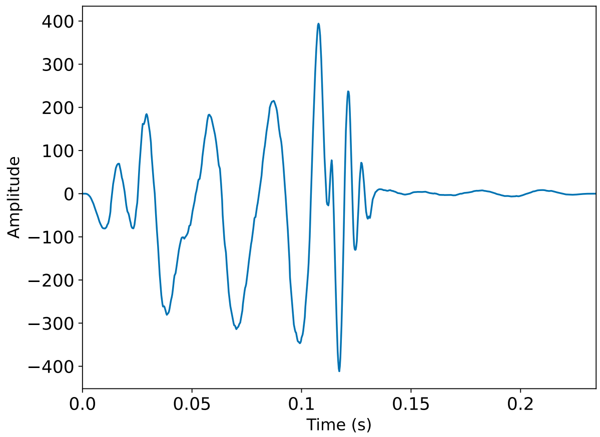

Figure 5.5: An example of an extreme false positive that is shared by the three classifiers. This glitch waveform resembles the waveform of an intermediate-mass black hole binary, where the inspiral phases are difficult to detect.

though the classifiers performed very differently on the signals, this does not seem to be the case for the glitches. As the glitches are shared between the datasets, the classifiers have likely learned similar features for these glitches. All of the found false positives strongly resemble signals, as is shown for one particular glitch example in Fig. 5.5.

The results from this investigation of the most extreme examples show that for the signals, the mass ratio and inclination are parameters of interest, as these parameters are higher on average for the false negatives. It is also retrieved that the duration plays a role. This mirrors the duration being a factor in matched filtering searches. In the case of the neural networks, it is possible the short duration signals do not exhibit enough of the features the neural networks have learned to make a positive classification. The misclassified glitches largely intersect, and the three classifiers show near-identical performance on the extreme examples. This likely means that even though the three classifiers do show different behaviours on the signals, the different training regimes has not effected the features learned to identify glitches. Based on the fact that the curriculum classifier performed well on the false negatives of both the aligned-spin and precessing-spin classifiers, one can argue that the curriculum classifier has successfully absorbed information that the aligned-spin classifier did not learn. This can be seen as support for the architecture not being fully robust to precession.

5.3.4 Parameter Sensitivity

The sensitivity of the classifiers to injection parameters can be investigated by relating the parameters to the classifier outputs. In this subsection this analysis is done first for the curriculum classifier, as it has the highest discriminating power, seemingly capturing the difference between the injected signals and glitches best.

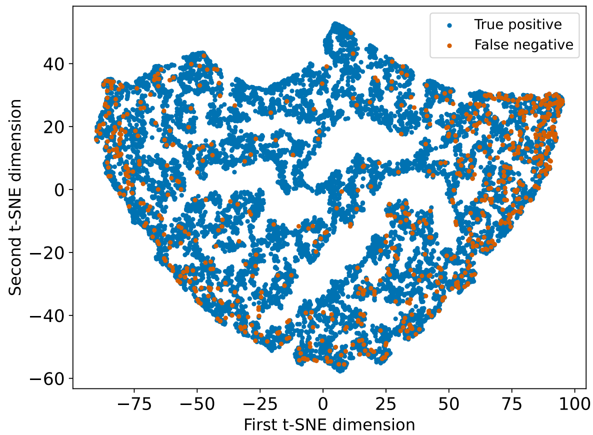

Figure 5.6: The true positives and false negatives of the curriculum classifier projected onto a plane using t-SNE with a perplexity value of 50. Different hyperparameters did not result in noticeable differences in the separation of the classes.

The obtained information can then be used for a statistical comparison for parameter sensitivity between the three classifiers as a test of robustness.

First, it is tested if separated clusters can be found in the class of signal examples using PCA and t-SNE. The existence of such clusters would give indication of sensitive regions for either the true positives or false negatives within the parameter space. Since every signal example that is assigned an output by the classifier was generated from a row in the injection table, the binary label prediction can be directly connected to the waveform parameters. The classifications can therefore be treated as data points residing in the parameter space, with each corresponding to either a true positive or false negative. This dataset was input to the t-SNE algorithm which visually maps the data down to a 2-dimensional space, as shown in Fig. 5.6. As can be seen from the figure, the output reveals no separated clusters of either class.

As an additional test, the same data was input to PCA, which captured nearly the full variance (99.6%) with two principal components. Here too however, the classes show no separation, even though the density of the false negatives does seem to correlate with the second principal component. This is shown in Fig. 5.7.

The results from both t-SNE and PCA suggest that no separated clusters of classifications exist, or that they can not be captured by these two algorithms. This reduces the current study of sensitivity to outlier analysis using individual statistics, rather than studying specific regions of the parameter space.

One way to search for outliers as well as to detect possible relations between

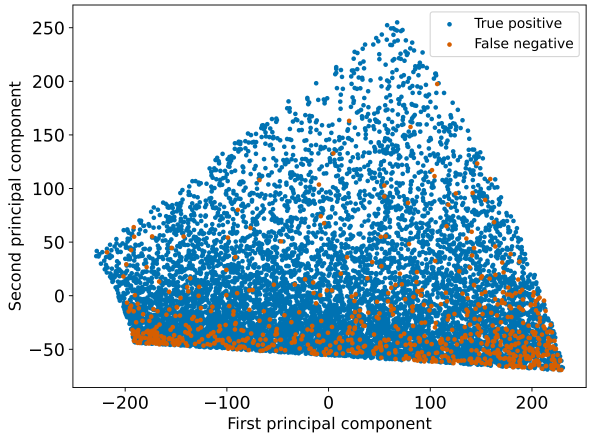

Figure 5.7: The true positives and false negatives of the curriculum classifier projected onto a plane using PCA. As for t-SNE, there is no separation between the classes.

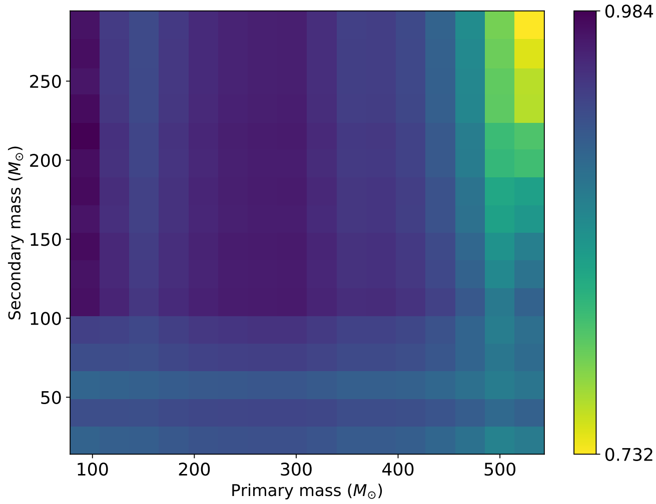

Figure 5.8: A heatmap of the outputs from the curriculum classifier on the grid of the primary and secondary masses. This heatmap reveals a sensitivity to higher total masses in the top right corner.

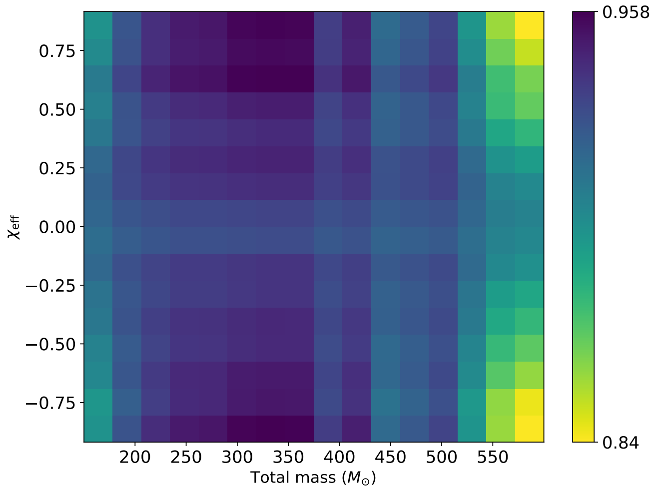

Figure 5.9: A heatmap of the outputs from the curriculum classifier on the grid of the total mass and χ_{eff}. It can be seen that there is a natural symmetry on the effective spin, and the classifier is proportionally sensitive to higher absolute values of this parameter.

parameters is to plot heatmaps, in which bins are created for the combinations of the values on the axes, and the average of the classifier outputs within these bins is shown. Creating heatmaps involving the inclination, the first parameter of interest found in the previous subsection, revealed no new patterns. For the primary and secondary masses, a heatmap is shown in Fig. 5.8. Although this figure shows that the curriculum classifier is on average correct and fairly confident (with a minimum cell average of approximately 0.73), uncertainty sets in in the highest region of total mass, or as the primary mass increases. This reflects the correlation of matched filtering sensitivity with the primary mass shown in [264], although the limitations of the curriculum classifier set in for total masses that are considerably higher. A direct comparison is made difficult by the difference in analysis, as [264] studies the probability of matched-filtering triggers being produced in these ranges.

Resuming the analysis on the total mass with the addition of the effective spin χ_{eff} (Subsec. 2.1.6), a similar heatmap is shown in Fig. 5.9. As expected, this heatmap reveals regions of higher uncertainty at high total mass, but also at the extreme ends of χ_{eff}. This is again in line with the findings of [264], and together these observations show that to a certain extent the curriculum classifier qualitatively behaves similar to how matched filtering was found to behave in earlier work, but with high confidence even at extreme parameter values. It is possible that convolutional neural networks to some degree learn the same filters as are used in matched filtering, as for some applications convolutional neural networks tend to repeatedly learn similar filters [265, 266]. This might suggest an interesting connection between the convolutional filters used in matched filtering, and the filters learned by the convolutional neural networks. Further investigation

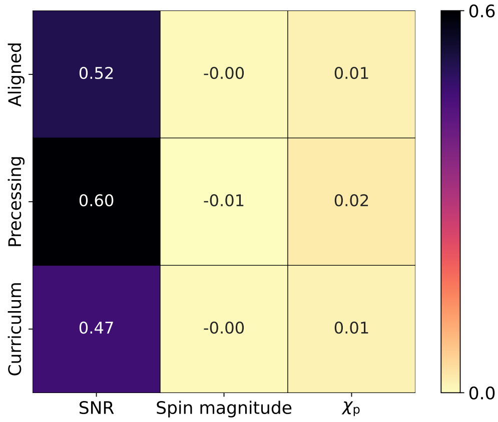

Figure 5.10: The values for the Spearman bivariate correlation coefficient with the classifiers on the vertical axis and quantities on the horizontal axis. This matrix shows that of the variables with the strongest correlation to the classifier outputs, the SNR is by far the best predictor.

is deferred to future work. Due to the horizontal symmetry in the heatmap, no further conclusions for χ_{eff} can be drawn from this plot.

In order to quantify by how much the outputs of the classifiers are influenced by the injection parameters and quantities derived thereof, bivariate correlation coefficients were computed for all three classifiers. In all cases the Spearman correlation coefficient yielded stronger evidence than the Pearson coefficient did, but with similar trends. The strongest correlations are shown in Fig. 5.10, clearly showing a moderate correlation of the classifier outputs with the SNR. This shows that the classifiers are mostly hindered by the visibility of the signals within the noise, as the correlation coefficients are only computed for signals, and therefore should never be assigned 0 by a perfect classifier. It is worthy of note that the correlation of χ_{p} with the output is higher for the precessing-spin classifier than it is for the aligned-spin and curriculum classifiers. This means the precessing-spin classifier is more likely to assign a higher output to examples that show more precession.

From the analysis in this subsection it appears that there is no separate region of the parameter space to which the classifiers are particularly sensitive, and that the binary black hole properties that trouble the classifiers most are akin to those hindering matched filtering searches. The values of the correlation coefficients suggest that beyond a certain point, obtaining correct classification simply becomes a problem of visibility, meaning there is not one specific parameter or statistic that can be appointed as a predictor for misclassification. In terms of robustness to precession, although not strongly correlating, the precessing-spin classifier does respond better to examples with a higher value of χ_{p} than the other classifiers do. This suggests that there are properties of precessing examples that the aligned-spin classifier and curriculum classifier have not fully learned.

Chapter 6

Accelerating Waveform Generation

Gravitational-wave analyses depend heavily on waveforms that model the evolution of compact binary coalescences as seen by observing detectors. In many cases these waveforms are given by waveform approximants, models that approximate the amplitude and phase of the waveform at a set of frequencies. Because of their omnipresence, improving the speed at which approximants can generate waveforms is crucial to accelerating the overall analysis of gravitational-wave detections. An optimisation algorithm is proposed that can select at which frequencies in the spectrum an approximant should compute the power of a waveform, and at which frequencies the power can be safely interpolated at a minor loss in accuracy. The algorithm used is an evolutionary algorithm modelled after the principle of natural selection, iterating frequency arrays that perform better at every iteration. As an application, the candidates proposed by the algorithm are used to reconstruct signal-to-noise ratios. It is shown that the `IMRPhenomXPHM` approximant [90] can be sped up by at least 30% at a loss of at most 2.87% on the drawn samples, measured by the accuracy of the reconstruction of signal-to-noise ratios. The behaviour of the algorithm as well as lower bounds on both speedup and error are explored, leading to a proposed proof of concept candidate that obtains a speedup of 46% with a maximum error of 0.5% on a sample of the parameter space used.

6.1 Introduction

Gravitational waves were first predicted well over a hundred years ago by Heaviside, Poincaré and Einstein, with Einstein famously remaining skeptical of their existence throughout his life. Einstein also believed that the amplitude of gravitational waves reaching earth would be so small that their detection could never be within human reach. Despite the accurate prediction of their magnitude, the LIGO-Virgo collaboration succeeded in making their first detection of a gravitational wave in 2015 [17]. Since, the LIGO-Virgo-KAGRA collaboration has confirmed detection of over 90 gravitational waves [42, 142–144]. Gravitational waves are now used to study black holes, neutron stars and other cosmological sources of gravitational waves.

In order to cover many possible source systems during a matched filtering search, templates for different systems are collected in template banks [135]. In turn, the templates populating this bank are generated by waveform approximants. Waveform approximants [88, 95, 267–269] serve to approximate the waveform of a system either in the time or frequency domain and are faster alternatives to the more precise waveform simulations from numerical relativity [80–82, 192, 270]. Approximants balance the need for both speed and accuracy, and therefore approximate a waveform to a reasonable accuracy [271] within a reasonable time. Other applications for approximants are in parameter estimation [139], where waveforms are generated during runtime, so that there too computational speed is crucial. If waveforms are computed during runtime this is called online computation, whereas it is called offline computation if this task is delegated to the pre-processing stage.

Since the Fourier transform provides a duality between the time domain and frequency domain, the domain of analysis is a choice. Sometimes it is preferential to work in the frequency domain, where a different perspective on data is allowed, or computations are easier to perform. This is the case for the methods presented in this chapter, where the frequency contents can be considered the building blocks of a waveform. A waveform approximant generating a waveform in the frequency domain will do so by solving sets of equations at individual frequencies in the spectrum [90]. Doing so is costly, notably when waveforms are generated online during the runtime of a procedure, but also in generally large undertakings such as template bank validation where percentual changes make a large absolute difference [272–275]. One way of reducing the cost of computation would be to evaluate a waveform on as small an array of frequencies as possible, resorting to alternative, faster procedures to determine the behaviour at the remaining frequencies. This is the premise of this chapter.

With this line of reasoning, a method that can speed up waveform approximants is calling the solver for a subset of the frequencies, and interpolating the values of the frequencies in the complement within the frequency spectrum. Finding where in the spectrum to solve and where to interpolate is then an optimisation problem, in particular one with a very large solution space that has size growing exponentially in the number of frequencies. One way to efficiently find a (local) optimum for this problem is through the use of evolutionary algorithms [179].

Evolutionary algorithms mimic natural evolution, where successive generations are expected to be better adapted to their environments. This principle translates to an algorithm that at each iteration tests a set of solutions, collectively called a generation, and then forms a new generation that contains better solutions to the problem. Evolutionary algorithms, sometimes also referred to as genetic algorithms, have been applied to gravitational-wave data analysis in other work [256, 276, 277].

Alternative methods with the same motive have been proposed, such as multi-band template interpolation [278] and adaptive frequency resolution [279]. However, both methods make assumptions on the frequency evolution of the waveform to reduce the number of frequencies at which the power is computed, and adaptive frequency resolution in particular was shown to be accurate mostly for relatively loud signals. Evolutionary algorithms require no assumptions to be made and are therefore more widely applicable, with no specific boundary on the loudness. Other related methods include the work done on reduced order quadratures [280–282] and meshfree approximations [283].

The application of evolutionary algorithms to the optimisation problem of waveform generation is studied in this chapter as a proof of concept, proposing an algorithm, and studying both the time speed-up and the coverage of the parameter space. Coverage here means the generalisibility of the frequency array output by the algorithm to samples taken from the parameter space, or on what subset of the parameter space the interpolation from the output frequency array gives accurate approximations.

In this chapter binary black holes are considered. If instead of a binary black hole a system containing a neutron star is considered, an additional parameter describing the tidal deformability of the star is added to this space [284]. This chapter will however restrict itself, and therefore the parameter space, to the case of black holes.

A waveform approximant defined on the parameter space $\mathcal{P}$ defined in Subsec. 2.1.6 can then be interpreted as a function:

$$A_{\mathbb{F}} : \mathcal{P} \to D(\mathbb{F}, \mathbb{C}) \tag{6.1}$$

that maps a given set of parameters into the set of discrete complex functions on the frequency spectrum $\mathbb{F} \subset \mathbb{R}$, which should not be confused with an algebraic field. Note that $D(\mathbb{F}, \mathbb{C})$ can alternatively be written as $\mathbb{C}^{\mathbb{F}}$, and that the notation $A_{\mathbb{F}}$ specifically includes the spectrum $\mathbb{F}$. Such a waveform approximant is typically the best possible option for the modeling of a system's evolution short of an analytical or numerical solution of the Einstein field equations with stress-energy-momentum tensor defined by the parameters of choice. This will be discussed in more detail shortly.

In this chapter `IMRPhenomXPHM` will be used for $A_{\mathbb{F}}$, as it is fast, commonly used online, and incorporates both precession and higher-order modes. Note however that the methods presented in this chapter are agnostic to this choice and are applicable to all approximants. As will be demonstrated in this chapter, obtaining A_F for a subset $F \subset \mathbb{F}$ and interpolating A_F to A^*_F on $\mathbb{F}$, where the asterisk indicates interpolation, is faster on a per-frequency basis than solving for all frequencies in $\mathbb{F}$, yet still accurate. The subset F will be referred to as the *frequency array*, and two coarse examples of such frequency arrays for a damped sinusoid are shown in Fig. 6.1. The convention of asterisks indicating linear interpolations will be

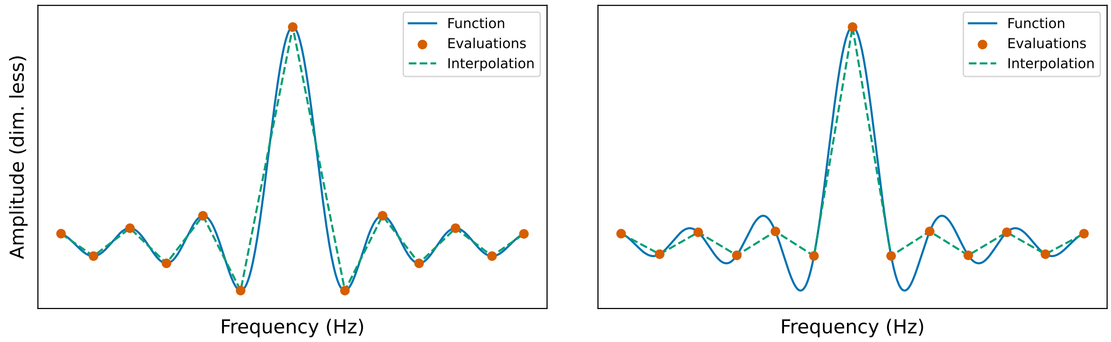

Figure 6.1: A damped sinusoid interpolated on two different frequency arrays. The left frequency array consists of the turning points of the function, whereas the right frequency array is uniformly spaced. This figure shows both how frequency arrays can be realised differently, as well as the importance of the choice of array. Although both contain the same number of frequencies, corresponding to the evaluations in the plots, the left interpolation is clearly a better approximation to the true function.

upheld. Notation will however be shortened, with the exact waveform $A_{\mathbb{F}}(\theta)$ being denoted by $h(\theta)$, and the approximate waveform $A^*_F(\theta)$ by $h^*_F(\theta)$. If the choices of θ and F are implicit, the notation can be further reduced to h and h^* respectively.

Since this speed-up is obtained through interpolation and not algorithmically, h and h^* can not be expected to be equal, and a loss in accuracy will be incurred. Two metrics will be used to measure this loss. The first is the residual sum of squares (RSS) for the real and imaginary parts of h and h^*. This metric is used by the evolutionary algorithm internally. The second metric is based on a widely used quantity, the signal-to-noise ratio (SNR).

6.2 Methodology

With the optimisation problem and solution method defined, this section discusses the methodology, split into three subsections. The pre-processing section discusses all considerations that preceed the introduction of the evolutionary algorithm to the problem. Following this, the hyperparameters that define the evolutionary algorithm for the established problem of finding the optimal frequency arrays are discussed in the second subsection, with the third subsection offering the performance measures used to evaluate the proposed solutions. Three such solutions will be considered.

In typical applications the frequency spectrum $\mathbb{F}$ is a uniformly spaced set, and in the tests presented here the set is defined by upper and lower limits of 0 and

2048 Hz respectively, with a sampling frequency of 4 kHz and a duration of 16 seconds.

6.2.1 Pre-processing

In order to obtain a faster rate of convergence for the evolutionary algorithm, reference events are used to populate initial frequency arrays during pre-processing. An event is defined here as a fixed set of parameters, and in the case of a catalogued candidate, the parameters used are retrieved from the parameter estimation data releases [42, 142–144]. Three such reference events are tested. First, GW150914 [17], as it was a landmark detection without additional complexities. Second and third are GW190412 [285] and GW190814 [218] since there is evidence for the contribution of higher-order modes in these events, as well as precession [286, 287]. For this proof of concept, only the plus-polarised components are used.

Starting from the parameters of a reference event, the parameters other than the chirp mass are fixed, and the chirp mass is varied by ± 5 solar masses to generate a number of different waveforms on $\mathbb{F}$. The choice to vary the chirp mass was made as it induces a large variation in the phase of the resultant waveforms [72]. Each of the waveforms are normalised in amplitude to $[-1, 1]$, and the turning points in the frequency domain are added to a set B called the *background* or prior. Every frequency array for this event will include this ever-present background of frequencies. This is done under the assumption that the evolutionary algorithm finding and including turning points is inevitable, since a piecewise linear interpolation between the turning points of a function generally gives a very well performing approximation to the full function. The inclusion of this prior information in pre-processing greatly reduces the work the evolutionary algorithm would have to perform in identifying these frequencies during a random walk. Choosing the right number of waveforms, or equivalently the permitted cardinality of B, is an example of bias-variance tradeoff [159]. For the tests described here, 5 waveforms were used, including the reference event.

Finally, all frequencies under 20 Hz were removed as candidates for inclusion in a genotype. The result of these pre-processing steps is one set of frequencies (the background) per event that will be present in all frequency arrays learned on the corresponding event, as well as one permissible set of frequencies to include in the genotypes for each.

6.2.2 Hyperparameters

Given the pre-processed data, the hyperparameters uniquely define the evolutionary algorithm up to stochasticity. In this subsection the different hyperparameters are treated, guided by the context wherein they arise. The reader is recommended to look ahead at Table 6.1 before reading this subsection in order

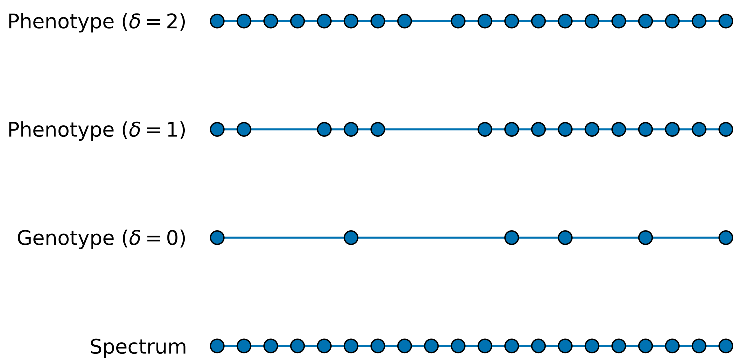

Figure 6.2: An example genotype and corresponding phenotypes for a spectrum of subsequent natural numbers with empty background. For $\delta = 0$ the genotype and phenotype coincide. With $\delta = 1$ nearby frequencies are added to the genotype to form the phenotype, and with $\delta = 2$ additional nearby frequencies are added. This figure shows the hyperparameter δ controls the confidence with which the algorithm includes frequencies in the solution phenotypes.

to get an overview of how this subsection is structured.

Recall that for a set S and a subset $T \subset S$, an indicator vector I_T is a binary vector indexed by the elements of S such that $I_T(s) = 1$ if and only if $s \in T$. Genotypes, as subsets of the frequency spectrum, are encoded in indicator vectors over $\mathbb{F}$, and a phenotype is constructed from such a genotype by adding the indicated frequencies to the frequencies in the background B. This map therefore realises a frequency array as a phenotype. Additionally, a parameter $\delta \in \mathbb{R}_{\geq 0}$ representing the confidence of the algorithm can be set: if a frequency f is included in the genotype, then the elements of the set:

$$\{f' \in \mathbb{F} : |f - f'| \leq \delta\} \tag{6.2}$$

will be included in the phenotype. Note that this means that with $\delta = 0$ only the frequency f itself is included in the phenotype, therefore reducing to the special case described above. The phenotype (or frequency array) construction process for different values of δ, with B set to the empty set for simplicity, is shown in Fig. 6.2.

For a genotype G (with corresponding phenotype F) and a reference event waveform h, the objective function is defined as:

$$O(G) = \sum_{f \in \mathbb{F} - F} |h(f) - h^*_F(f)| + w \cdot \#F, \tag{6.3}$$

where $|\cdot|$ denotes the modulus and the hash denotes set cardinality. The sum of moduli amounts to the residual sum of squares of the exact and approximate waveforms for both the real and complex parts, as can be seen from expanding the modulus for a single frequency. This sum can therefore also be expressed as:

$$\mathrm{RSS}(\mathrm{Re}[h], \mathrm{Re}[h^*_F]) + \mathrm{RSS}(\mathrm{Im}[h], \mathrm{Im}[h^*_F]). \tag{6.4}$$

An additive penalty, weighed by a variable w, is added for the phenotype size. One should take note of the range for the summation in the objective function defined in Eq. 6.3: for all frequencies in F the exact waveform coincides with the approximate waveform, meaning these terms reduce to zero, and therefore do not have to be computed.

The complexity of the solution space, and therefore the need for optimisation methods like evolutionary algorithms, is demonstrated through the objective function. A first approach to optimisation could be by way of a greedy algorithm [136], where at every iteration the frequency f' corresponding to the largest residual, or equivalently the largest term in the sum, is added to the phenotype. However, due to the addition of this frequency, the approximate waveform is changed to

$$h^*_{F\cup\{f'\}}, \tag{6.5}$$

and the terms of the sum in the objective function will have to be recomputed. At the same time the cross-section of the (very high-dimensional) solution space containing solutions without the frequency f' is pruned. There is no guarantee that the frequency f' that is now included is part of the optimal solution as this objective function has no optimal substructure [136], and no choice of frequency is ever reconsidered. It follows that the greedy algorithm would now be barred from finding a global optimum, whereas an evolutionary algorithm could rectify this misstep.

If a phenotype frequency array is used for the computation of the SNR, the error and rate of convergence are quantifiable, and minimising the sum of moduli in the objective function is equivalent to minimisation of the error term. Assume D is an arbitrary but fixed strain of data, and h is a waveform. It then follows from the linearity of the inner product that the error $E(h)$ for this waveform is given by:

$$E(h) = \langle h, D\rangle - \langle h^*, D\rangle = \langle h - h^*, D\rangle. \tag{6.6}$$

Since taking the limit of an inner product is interchangeable with taking the limit of an argument, and the continuous inner product will tend to 0 if either argument tends to 0, it is given that $E(h) \to 0$ if $h - h^* \to 0$. The latter is exactly the case if $|h - h^*_F| \to 0$ in the L^1-norm.

Once the objective function has been used to evaluate the fitness of the genotypes in a generation, a new generation of offspring can be formed. The parent selection function P is implemented according to fitness proportionate selection with elitism [34, 288]. Elitism refers to the carrying over of the top percentage of genotypes (by fitness ranking) to the next generation, and fitness proportionate selection means that the higher a genotype is ranked, the higher the probability will be that it is drawn as a parent by P. The probability measure so constructed on the genotypes sorted by ascending fitness is given by $\mathbb{P}(i) = i \cdot p$ with

$$p = \frac{1}{\sum_{j=1}^{g} j} = \frac{2}{g(g+1)} \tag{6.7}$$

for the candidate ranked at position i, and g the fixed population size of a generation. The use of this measure ensures that the lowest ranked genotype (at $i = 1$) has the lowest probability of being chosen as a parent, with every higher ranked genotype having a proportionally higher probability of being selected. The second equality in Eq. 6.7 follows by substituting in Gauss' summation formula. Note too the importance of the generation size, as for larger values the probabilities will be more evenly distributed over the candidates.

Finally, the recombination function R is defined as the sum of the two genotypes, assuming the genotypes to be a boolean algebra. This means that a frequency will be present in the offspring if and only if that frequency is present in at least one of the parent genotypes.

Due to a symmetry argument applied to the phenotype construction process, the algorithm can be ran both forwards and backwards. In the forwards case, the phenotype grows in size towards the full frequency spectrum as frequencies are added in each iteration of the algorithm. In the backwards case the frequency spectrum is reduced in size towards the background by removing frequencies in each iteration. As a rule of thumb the forwards case is better suited to finding smaller frequency arrays, whereas the backwards case is better suited to finding frequency arrays that result in a low error. This is a result of the stopping conditions typically being met before the frequency array reaches half the size of the frequency spectrum. Note that the parameter that dictates the direction stands in direct relation to the penalty w included in Eq. 6.3. In both cases the algorithm needs to be incentivised to move through the solution space by the penalty. In for instance the backwards direction, if w were to be set to 0, there would be no reason to remove any frequency from the phenotype, as doing so would only lead to a loss in accuracy. Setting $w > 0$, the removal of this same frequency might result in an increased fitness, as long as the loss in accuracy is less than the penalty w for including the frequency.

The evolutionary algorithm that has now been established can only be run once an initial generation is given. Initialisation of this first generation is done at random through a uniform distribution on the set of allowed frequencies. A generation of genotypes is formed by drawing a sample of frequencies of size s from this distribution, repeating this until g such genotypes have been drawn to form the full first generation. This marks the frequency sample size s as a hyperparameter. Note that the number of possible genotypes that can be drawn is given by a binomial coefficient, and due to the factorial in this coefficient this number will be very large. In order to explore a large section of the solution space, the sample size s

Parameter	**Type**	**Description**	**Value**
δ	Float	Confidence	0
w	Float	Penalty weight	8
g	Int	Population size	6000
d	Boolean	Algorithm direction	`True`
s	Int	Population sample size	6
r	Int	Number of generations	30

Table 6.1: The collection of hyperparameters that defines the evolutionary algorithm. The value column contains the values for these parameters used for the realisation of the algorithm.

needs to be chosen proportional to both g and the maximum runtime r of the algorithm expressed in the number of generations. In the k-th generation, a maximum in the order of magnitude of $s \cdot 2^k$ frequencies are included in the genotype. Using this information an appropriate combination of values for s, g and r can be set.

Recapitulating this subsection, besides the initial choices of background size during pre-processing and the probability measure for the parent selection function, the algorithm is defined by the hyperparameters (δ, w, g, d, s, r) given in Table 6.1. The values in this table represent the hyperparameters used in this proof of concept, with `True` for the direction d meaning the algorithm was ran backwards. This choice was made to prioritise accuracy, which also inspired the choice of δ. The runtime r was set to well above the observed limit of stagnation in trial runs, with the population size g and frequency sample size s empirically determined following the previous discussion. The penalty weight w is set to 8, but any value to counterbalance the increasing modulus in the objective function as the evolutionary algorithm progresses and frequencies are removed would have sufficed. These choices of hyperparameters consistently drive the evolutionary algorithm to converge to similar and well-performing frequency array phenotypes on repeated runs, suggesting the optimisation does not get stuck in easily avoidable local minima. The runtime for the algorithm (including pre-processing) on an Intel Gold 6154 CPU is in the order of magnitudes of minutes, bound by the longest recorded run.

6.2.3 Performance measures

In the remainder of this section two metrics to measure performance will be introduced. The first is a fractional measure of the speedup obtained by using a frequency array phenotype output by the evolutionary algorithm. The second measures the loss in accuracy for a frequency array, where the exact waveform was injected into LIGO Livingston detector noise before the SNR was computed by matched filtering using both the exact and approximate waveforms as templates. The percentual difference in SNR is then used as a measure of loss.

For the definition of the speedup metric, let t_a be the average time in seconds elapsed by the approximant per frequency, and t_i the average time required to interpolate at a frequency in seconds. With $\rho := \#F/\#\mathbb{F}$ the relative size of the frequency array, the speedup of F can be computed as:

$$S[F] = \frac{\#\mathbb{F} \cdot t_a}{\#\mathbb{F} \cdot [\rho t_a + (1-\rho)t_i]}, \tag{6.8}$$

since the numerator specifies the time if the waveform is computed at every frequency, and the denominator specifies the time if the waveform is computed at a fraction ρ of the frequencies, with interpolation at a fraction $(1-\rho)$ of the frequencies. This expression reduces to:

$$S[F] = \frac{1}{(1-\rho) \cdot t_i/t_a + \rho}. \tag{6.9}$$

Note that this quantity is dimensionless, measuring the speedup as a ratio in an isolated section of a procedure where the waveform is generated, without depending on the machine. In order to obtain accurate values for t_a and t_i, 100 waveforms were generated on the full spectrum using `Bilby` [155], and the same 100 waveforms were interpolated on the frequency array phenotype. For every frequency array the parameters of the corresponding reference event were used without loss of generality since only time is measured, with the primary mass perturbed by ± 1 at every iteration to ensure no values were cached. The times per frequency were then averaged out for each event and type, repeating and averaging this process 100 times to remove possible inconsistencies.

Finding the neighbourhood in the parameter space $\mathcal{P}$ where the frequency arrays are effective is important, as the reference event input to the algorithm during pre-processing adds a level of specificity. Nearby the reference events the frequency arrays are likely to allow a more accurate reconstruction of the SNR. In order to analyse the generalisibility, $10,000$ points are sampled from the default precessing binary black hole prior in `Bilby` version 2.2.2. At every sampled set of parameters, the exact plus-polarisation of the waveform is generated and injected into simulated noise generated using the LIGO Livingston PSD included in `Bilby` version 2.2.2. The SNR of both the exact plus-polarisation, hereafter also called the exact SNR, and the approximate plus-polarisation, hereafter called the approximate SNR or SNR*, is computed against this strain. The percentual error for the frequency array F is then calculated as:

$$E_{\%}[F] = \frac{|\mathrm{SNR} - \mathrm{SNR}^*|}{\mathrm{SNR}} \cdot 100. \tag{6.10}$$

This metric is computed as a percentage to account for the different values of the exact SNR. A given absolute error might be negligible for large SNR values,

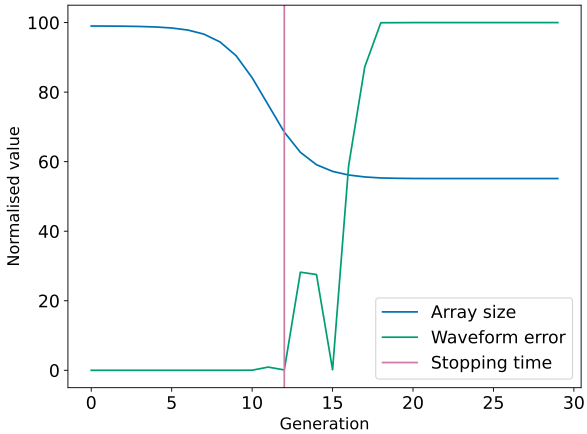

Figure 6.3: A visualisation of the two parts that make up the objective function for GW150914, along with the stopping time for this event (generation 12). The waveform error is normalised to $[0, 100]$ so that the graphs can be compared at the same scale.

whereas it might be detrimental to a low value SNR reconstruction. Computing the error as a percentage removes this consideration. A threshold of 0.01 is set for $E_{\%}[F]$, and a parameter sample is said to be rejected if the percentual error either attains or exceeds this threshold. Considering the uncertainties in SNR values in the GWTC-3 catalog [42] this threshold is particularly strict. This treshold is used as an aid to quantise results and is inconsequential to their objective interpretation.

6.3 Results

In this section the results from running the evolutionary algorithm with the three reference events GW150914, GW190412 and GW190814 are presented. The three frequency arrays output by the evolutionary algorithm are analysed in terms of speedup, accuracy and their generalisibility before selecting one of the three based on how well these three qualities are balanced. This frequency array is used for additional tests.

The results for the reference events are shown in Fig. 6.3, 6.4 and 6.5 respectively, where both constituents of the objective function are shown in a visually separable way. The waveform error in these figures refers to the residual sum of squares, normalised to $[0, 100]$ so it is more easily compared to the percentual array size at this same scale. Normalisation here means the minimum achieved waveform error is mapped to 0, and the maximum achieved waveform error is mapped to 100. This means that the visualised waveform errors give little qualitative information and can not be compared across figures, because the original scales are not shown. The benefit however is that this allows for a visual rule to select the

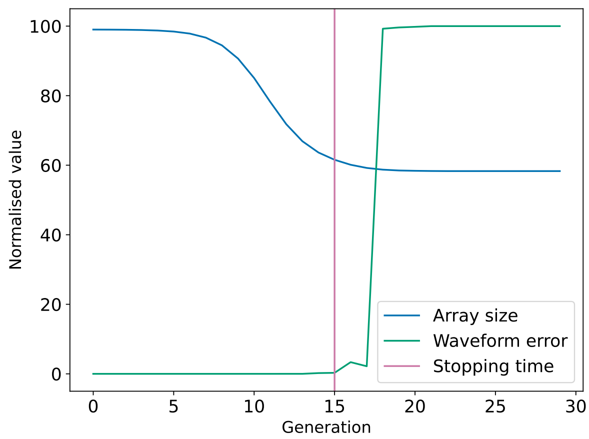

Figure 6.4: A visualisation of the two parts that make up the objective function for GW190412, along with the stopping time for this event (generation 15). The waveform error is normalised to $[0, 100]$ so that the graphs can be compared at the same scale.

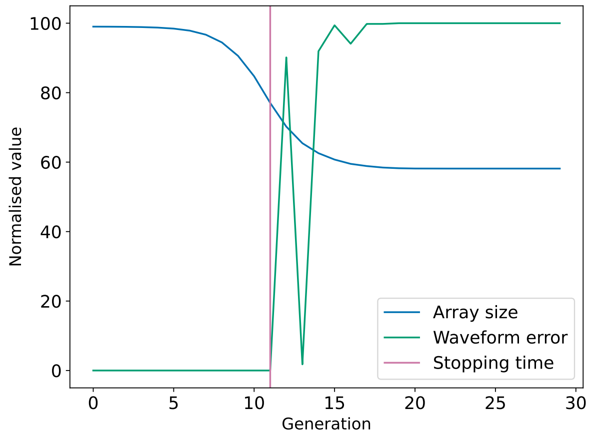

Figure 6.5: A visualisation of the two parts that make up the objective function for GW190814, along with the stopping time for this event (generation 11). The waveform error is normalised to $[0, 100]$ so that the graphs can be compared at the same scale.

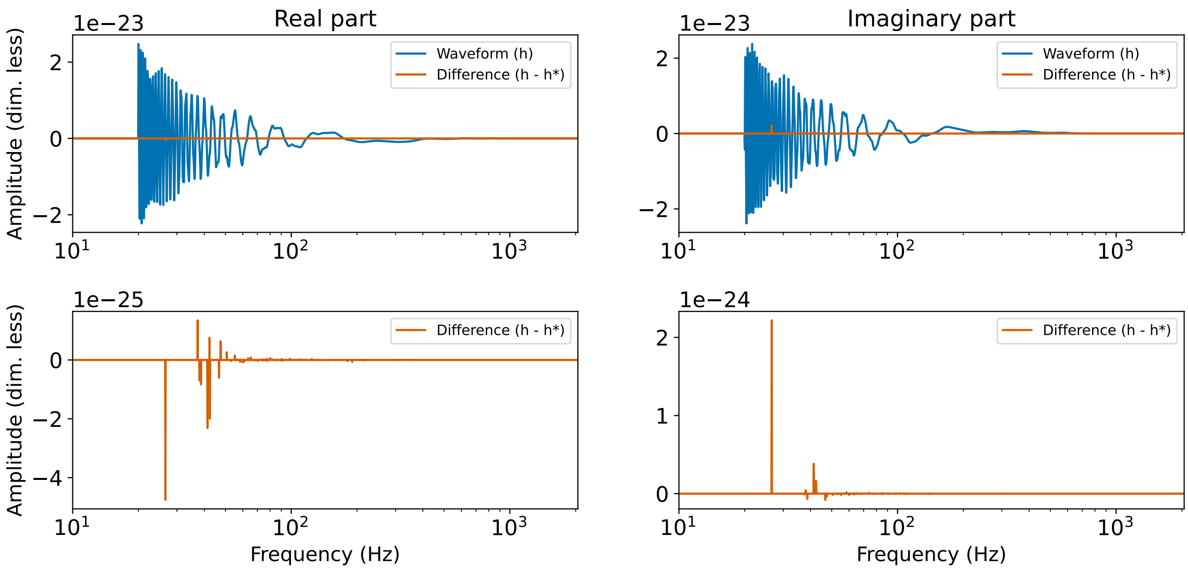

Figure 6.6: The real and imaginary parts of the plus-polarised component of GW190412 (h) and the difference with the approximation ($h - h^*$) computed on the frequency array proposed in the last generation (29) of the algorithm. The bottom two panels show the differences on a smaller scale, and the horizontal axes are in log scale.

stopping time, and one that is in fact automatable: halting at the first sign of a significant jump in error. At this time the evolutionary algorithm is stopped and the frequency array with the highest fitness score is chosen as the solution.

At first glance, the non-monotonicity of the waveform error might seem conflicting with the assumption that every subsequent generation contains better adapted solutions. However, one should keep in mind that this statement is made with respect to the objective function, where the waveform error is combined with a size penalty, and that the objective function will show monotonicity. Overall, the fact that the waveform error is only relative is underlined by the difference between the exact and approximate waveform visualised in Fig. 6.6. Despite the dramatic jump in waveform error seen in Fig. 6.4, the difference $h - h^*$ for the frequency array proposed in the last generation (where the error is maximised) barely registers at the scale of the waveform, being at least an order of magnitude smaller. These magnitudes can be considered as upper bounds for the waveform errors over the runtime of the evolutionary algorithm.

In the current implementation, the error is more pronounced in the phase than it is in amplitude. Depending on the application, the objective function can be modified to either balance this error, or favor one of the two.

The typical development of the proposed frequency array for GW150914 over the

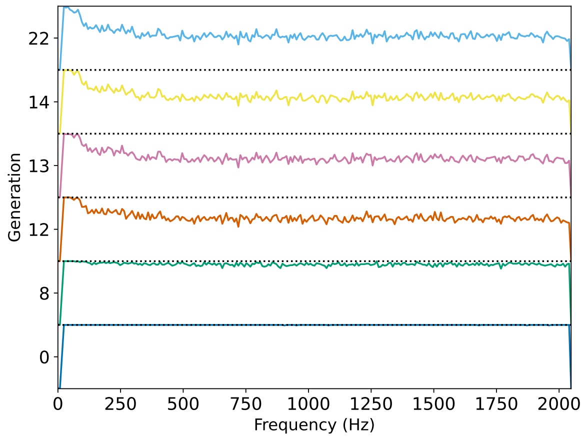

Figure 6.7: The evolution of the proposed frequency arrays for GW150914 visualised for a selection of generations. The densities are computed by binning the frequency spectrum, counting frequencies, and normalising to the maximum count. The evolution can be seen to follow the trend in the array size seen in Fig. 6.3, and shows what frequencies are removed during the runtime of the evolutionary algorithm. Notably, the frequencies between 0 and 250 Hz persist the longest.

generations is visualised in Fig. 6.7. The generations shown were selected to equally space the process before stagnation occurs at generation 22, and includes the ultimately selected candidate from generation 12. Both the stagnation and jump in waveform error that underlies this choice of candidate can be seen in Fig. 6.3. At every generation, a distribution function over the frequency spectrum is shown. This distribution is computed by dividing the spectrum into bins, counting the frequencies present in each bin, and normalising the counts to the maximum count over all bins. Mirroring the development of the array size in Fig. 6.3, few frequencies are removed from the early candidates, with only a minor visible move downwards from the dotted line seen in generation 8. Change accelerates between generations 8 and 12. Note that from the peaks and valleys that persist from this generation onwards, the algorithm has decided at an early stage which frequency ranges are most important. The lower end of the frequencies are kept consistently, although the dropoff towards 250 Hz becomes steeper during the later generations. Interpreting this figure gives an indication of how the evolutionary algorithm has traversed the search space, favoring the lower range of the spectrum, but always locally to the extent that above 250 Hz, the densities look like a diffused horizontal line. This means the evolutionary algorithm does not favor any band of the spectrum as strongly as it does $[0, 250]$ Hz.

The speed-up of the frequency arrays are shown in Table 6.2, showing the input variables and values for $S[F]$ as defined in Eq. 6.9. A detail that should be noted is that there is a natural inverse relation between the relative array size ρ

	ρ	t_a	t_i	**S**
GW150914	0.69	$5.52 \cdot 10^{-7}$	$6.95 \cdot 10^{-12}$	1.46
GW190412	0.62	$8.48 \cdot 10^{-7}$	$7.66 \cdot 10^{-12}$	1.62
GW190814	0.77	$1.71 \cdot 10^{-6}$	$5.90 \cdot 10^{-12}$	1.30

Table 6.2: Measurements of the values required for the computation of the speedup $S[F]$ given in Eq. 6.9, along with the computed speedup. These values are ρ, the size of the frequency array relative to the spectrum, t_a, the average time in seconds spent computing the waveform at a frequency by the approximant and t_i, the average time in seconds spent to interpolate the waveform at a frequency. All values are rounded to two decimals for readability.

and speedup S. The more frequencies are omitted from the frequency array, the higher the speedup. However, the values in the time columns differ to an extent. The differences in the t_i column can be explained by the density of the frequencies included in the array. The time measurements given in this column correspond to the average time elapsed interpolating at a frequency that is not included in the array. If many subsequent frequencies are omitted, the value at these frequencies can be interpolated at once from a single linear segment, whereas the same number of values for non-subsequent frequencies need to be interpolated using several such segments. The computation times scales linearly with the number of segments that are computed and will therefore be higher in the latter scenario. For the t_a column, the time difference between frequency arrays is likely to stem from the additional effects added to the waveforms, such as twisting up.

As a note on data compression and memory usage, before interpolating to the full frequency spectrum, one can opt to write the waveform to disk, indexed by the frequency array. By choosing to interpolate to the full spectrum later, waveforms can be stored in a compressed format, and the compression factor will equal ρ.

The results of the generalisibility tests to determine the regions of the parameter space $\mathcal{P}$ where the frequency arrays are effective are shown in Table 6.3. For each of the 15 parameters (minus the geocent time) the sampled ranges are shown, as well as the ranges that were recovered for the frequency arrays. Recovered here means that the parameter samples from these ranges were accepted by their error being below the given threshold, with $E_{\%}[F] < 0.01$. For most parameters, the full ranges were recovered, save for the luminosity distance and chirp mass. In the case of the luminosity distance, further inspection reveals that the lower ranges were simply not reached in the parameter sample drawn. The same holds for the chirp masses due to the way `Bilby` samples the chirp mass through uniform sampling in the component masses. If this is manually changed the results might vary slightly.

Additional statistics on the generalisibility tests are shown in Table 6.4. A number

	Range	**GW150914**	**GW190412**	**GW190814**
Mass ratio	$[0.125, 1]$	$[0.125, 1]$	$[0.125, 1]$	$[0.126, 1]$
Chirp mass ($M_\odot$)	$[25, 100]$	$[25.005, 86.922]$	$[25.007, 86.855]$	$[25, 86.509]$
Luminosity distance (Mpc)	$[100, 5000]$	$[101.091, 4999.683]$	$[170.862, 4999.658]$	$[128.847, 4999.707]$
Declination (Rad)	$[-1.571, 1.571]$	$[-1.552, 1.557]$	$[-1.553, 1.553]$	$[-1.545, 1.559]$
Right ascension (Rad)	$[0, 6.283]$	$[0, 6.282]$	$[0, 6.281]$	$[0, 6.283]$
Orbital inclination (Rad)	$[0, 3.142]$	$[0.003, 3.114]$	$[0.005, 3.138]$	$[0.026, 3.121]$
Polarisation angle (Rad)	$[0, 3.142]$	$[0, 3.141]$	$[0, 3.141]$	$[0.001, 3.141]$
Coalescence phase (Rad)	$[0, 6.283]$	$[0.001, 6.283]$	$[0, 6.283]$	$[0, 6.282]$
Spin of primary mass	$[0, 0.99]$	$[0, 0.99]$	$[0, 0.99]$	$[0, 0.99]$
Spin of secondary mass	$[0, 0.99]$	$[0, 0.99]$	$[0, 0.99]$	$[0, 0.99]$
Tilt of primary mass (Rad)	$[0, 3.142]$	$[0.014, 3.136]$	$[0.023, 3.126]$	$[0.018, 3.1]$
Tilt of secondary mass (Rad)	$[0, 3.142]$	$[0.004, 3.13]$	$[0.026, 3.127]$	$[0.017, 3.125]$
Precession angle (Rad)	$[0, 6.283]$	$[0, 6.283]$	$[0.003, 6.283]$	$[0, 6.283]$
Azimuthal angle (Rad)	$[0, 6.283]$	$[0, 6.282]$	$[0, 6.282]$	$[0, 6.283]$

Table 6.3: The ranges for each of the variables in the parameter space and the recovered ranges per frequency array, in rounded numerical format. If no unit is included, the quantity is dimensionless. Note that the geocent time is not included, as it is inconsequential to this test. All parameter ranges save for the luminosity distance and chirp mass were recovered, with both of these not being fully covered by the drawn sample of parameters.

	Minimum	**Maximum**	**Rejected**
GW150914	2.69e-10	0.50	909
GW190412	1.12e-07	2.87	902
GW190814	1.47e-09	0.21	203

Table 6.4: The extrema for the percentual errors $E_{\%}[F]$ on the parameter sample, along with the number of parameter samples rejected, as for these rejected samples $E_{\%}[F] \geq 0.01$.

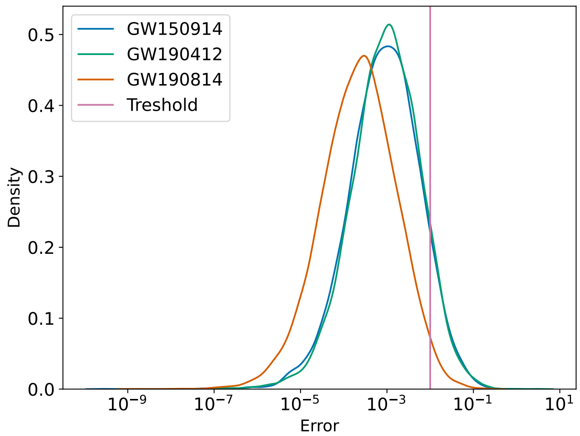

Figure 6.8: Densities of the percentual error in SNR $E_{\%}[F]$ for the full parameter samples, along with the threshold for acceptance. The horizontal axis is in log scale.

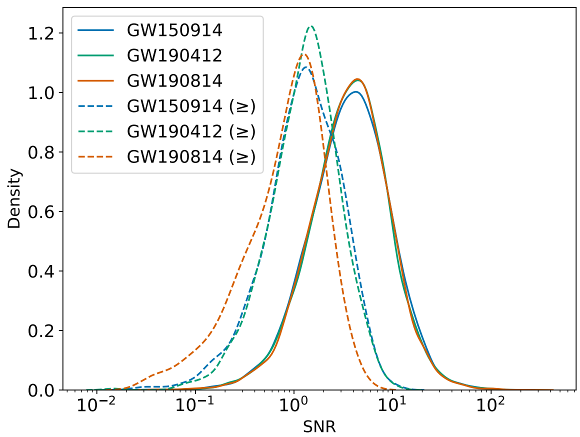

Figure 6.9: The exact SNR distributions for the full parameter samples (solid) and the rejected parameter sample points that exceed ($\geq$) the threshold of 0.01 set on $E_{\%}[F]$ (dashed). The horizontal axis is in log scale.

of parameter sample points were rejected. However, the errors for these samples were evaluated at a strict error threshold. Keeping in mind the generalisibility, these results can be considered remarkable, as the error is bounded above by 2.87%. The distribution of the errors is further visualised in Fig. 6.8. This plot shows that for all three frequency arrays the percentual errors in SNR are mostly concentrated below the error threshold of $10^{-2} = 0.01$.

Investigating the cause for the rejection of parameter samples, it becomes apparent that nearly all rejected samples have a low exact SNR value for h. This is visualised in Fig. 6.9, showing the exact SNR distributions for the rejected pa-

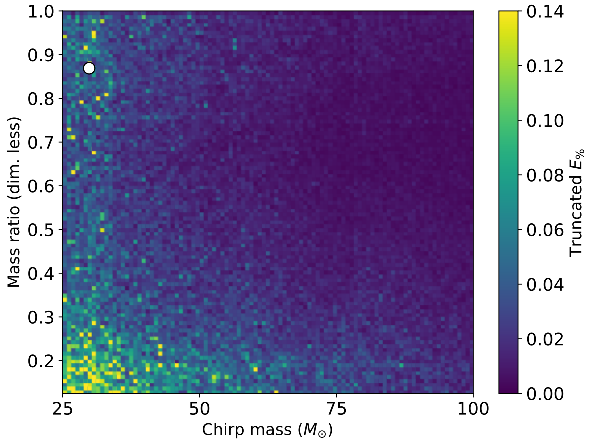

Figure 6.10: A heatmap for the error in SNR $E_\%$ based on partial parameter samples when using the proposed frequency array obtained using GW150914 as the reference event. A partial parameter sample means that a sample was drawn from the parameter space without sampling the chirp mass or mass ratio. For every such sample, the chirp mass and mass ratio were varied within the ranges of $[25, 100]$ and $[0.125, 1]$ respectively. The average errors for ten such samples are shown in this figure, truncated at 0.14% for increased visibility. The mass ratio and chirp mass of GW150914 are indicated by a white marker.

rameter samples versus the exact SNR distributions for the full samples. It can be seen from this figure that for the rejected parameter samples, the SNR is on average an order of magnitude smaller than that of the full parameter sample, and moreover that the distributions are left-tail heavy. In order to verify the strength of the SNR as a predictor for rejection, logistic regression [159] of the binary acceptance class onto the waveform parameters was employed. The sampled parameters for the three candidates were stacked, with the luminosity distance replaced by the SNR. The removal of the luminosity distance is a requirement so that the multicollinearity assumption of logistic regression is not violated, as luminosity distance correlates strongly with SNR. Despite the SNR scaling with several other parameters, the removal of the luminosity distance should sufficiently reduce the correlation. A binary labelling was created for this dataset by assigning a dependent variable 1 to the accepted parameter samples and 0 to the rejected samples. Fitting a logistic regression model, the SNR was found to be the predictor with the highest coefficient by far, with a value of $\beta_{\text{SNR}} = 16.5$. This coefficient was followed by the chirp mass with a value of $\beta_{\text{Chirp}} = 3.74$, meaning the SNR is a significantly stronger predictor for the rejection of a parameter sample according to the regression model. As with accuracy, the generalisibility too is bounded from below by the results presented.

Based on the analysis in this section, the frequency array obtained by using GW150914 as the reference event is proposed as a proof of concept, as it ob-

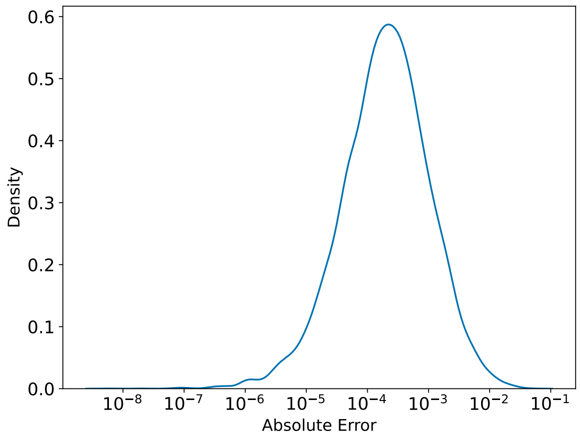

Figure 6.11: The distribution of the absolute error $|\mathcal{L}_m - \mathcal{L}^*_m|$ between the exact log likelihood maximum $\mathcal{L}_m$ computed using the exact waveform and the approximate log likelihood maximum $\mathcal{L}^*_m$ computed using the approximate waveform. This distribution was constructed on a sample of size $10,000$. The horizontal axis is in log scale.

tains a speed-up in computational time of 46% at a maximum error of 0.5% in SNR. A heatmap for the generalisibility of this frequency array is shown in Fig. 6.10. This figure was created by drawing 10 partial parameter samples, meaning every parameter except for the mass ratio and chirp mass were sampled, and then varying the mass ratio and chirp mass before computing the percentual error in SNR $E_{\%}$ for every such combination. A point in the plane then corresponds to the average error of the 10 parameter samples for that combination of the chirp mass and mass ratio, and the full figure shows $100,000$ such parameter samples in total. It can be seen from the figure that especially in the band of the chirp mass where the event belongs to, moving away from the mass ratio of the event will increase the error.

As an additional test for the applicability of the proof of concept, the maxima of the log likelihoods [138] that quantify the likelihood of waveforms being present in detector strains were approximated using the frequency array. Using the inner product defined in Eq. 3.7, the log likelihood can be defined as:

$$\mathcal{L}(h, D) = -\frac{1}{2}\langle D - h, D - h\rangle \tag{6.11}$$

for h the waveform projected onto the detector frame and D the strain. Setting $D := N + h'$ for N noise, the maximum of the log likelihood should be attained at the injected waveform h'. Analogous to the SNR reconstruction, the exact log likelihood maximum $\mathcal{L}(h, N + h)$, or $\mathcal{L}_m$ for shorthand, and approximate log likelihood maximum $\mathcal{L}(h^*, N + h)$, or $\mathcal{L}^*_m$, were computed for a parameter sample of size $10,000$. Here, however, both polarisations are included. Note that these errors do not need to be converted to percentages as was done for the errors in SNR, since the variance in the values will be significantly smaller. The absolute error distribution is shown in Fig. 6.11 and shows promising results. The reader should be aware however that the evolutionary algorithm was not finetuned to this computation, suggesting that a higher accuracy can likely be obtained. Specifically, the objective function (that defines the optimisation problem) should be adapted to the log likelihood function, and the hyperparameters should be rebalanced.

Further study in general might reveal marginally better performing frequency arrays, or frequency arrays better suited to specific sections of the parameter space. It is however expected that the improvement from frequency subsampling and linear interpolation represented by the values in this section is near-optimal. Implementing the naive frequency subsampling scheme of interpolating every second frequency (setting the relative array size at $\rho = 0.5$) for GW150914, the speedup S evaluates to 1.99, but the minimum error of $E_{\%}$ blows up by a factor with order of magnitude 10^4, in the process causing a rejection for 5972 parameter samples.

It follows from the results in this section that as expected, for a given threshold, smaller frequency arrays are effective on smaller neighbourhoods of the parameter space $\mathcal{P}$. Under the assumption of continuity, parameter instances nearby the reference event parameters will give better performance because of their similarity. If evolutionary algorithms are to be used to speed up computations on $\mathcal{P}$ with a given maximum allowable error, a single frequency array might not suffice. One option would then be to assign different frequency arrays to different regions of $\mathcal{P}$. In order to do this efficiently, the effective neighbourhood of an array will have to be computed. A first proposition could be akin to the generalisability test used in this chapter. After setting a treshold, recover the ranges in which all parameter samples were accepted, and take the convex hull of these ranges in $\mathcal{P}$. There exist efficient algorithms for the computation of convex hulls in higher dimensions under the assumption of flat space [289] that enable this procedure. Finding an optimal coverage of $\mathcal{P}$ by the resulting polytopes, however, reduces to an instance of the sphere packing problem through a topological argument, which depending on specifics is known to be NP-complete [290]. This obstacle can be overcome at the cost of efficiency if overlap of the polytopes is allowed and a global optimum is not required. Furthermore, for sufficiently large polytopes, the largest embeddable hypercube can be determined and used as an approximation to the polytope. These cubes in turn can be used to tile $\mathcal{P}$ more easily due to their shape. If such a tiling of $\mathcal{P}$ is generated the speedups demonstrated in this section can serve as lower bounds for the speedups achievable on the entirety of $\mathcal{P}$.

Chapter 7

Conclusions

7.1 Cosmic String Searches

A deep-learning model that can distinguish between cosmic string cusp signals and blip glitches with significant accuracy was designed and analysed. Given that matched-filter searches for short-duration gravitational-wave signals are heavily hindered by short transient glitches such as blip glitches, the exploration of this task is important for both current and future searches. In this work, both populations were scaled to follow the same SNR distribution, meaning loudness was removed from the equation. With remarkable results for the accuracy (79%) and true positive rate (76%) in particular, it has been shown that deep learning is a viable candidate for use in cosmic string searches. Moreover, due to the classification speed of 10 milliseconds per three data streams of 8 seconds, the deep-learning model is fast enough to run as part of a real-time detection pipeline.

On a dataset consisting of injected signals and injected glitches, the deep-learning model was shown to outperform matched filtering at the task of distinguishing strains including signals from strains including glitches, winning mostly on the volume of false positives (as can be seen from the model's slow increase in true positive rate in Fig. 4.16). This demonstrates that the deep-learning model is significantly better at rejecting glitches.

The behaviour of neural networks is notoriously difficult to understand, earning them the name of black-box models. As evidenced by their proven effectiveness, however, these black boxes hide valuable information. The hidden representations within the deep-learning model were interpreted through the application of three interpretability methods. The first of these is the method of waveform surgery, introduced in Ch. 4, where parts of a waveform are removed to study the effects on the classification of such a procedure. The second method is a routine developed in Ch. 4 for the visualisation of convolutional filter activations for one-dimensional time series. The third is principal component analysis. These interpretations have resulted in several observations that may prove useful in future work. The glitch surgery procedures demonstrate the possibility of dividing waveforms into sections and show that models can be sensitive to changes within these sections. Surgery procedures can therefore be used to study the importance of distinct sections of waveforms to their classification. By considering a comparison between waveforms based on their sections, the complexity of signal discrimination can be reduced, therefore potentially reducing the difficulty of the task. Through the study of the convolutional filter activations, these latter values can be connected to the classes in the confusion matrix. Such studies may aid in making informed choices for convolutional filters, or more generally in neural network design. Lastly, principal component analysis applied to the throughput of the second to last dense layer of the network has enabled the study of class separability and the hidden procedure reducing the internal representation of the deep-learning model to output probabilities.

In the process of interpreting the deep-learning model, the morphological differences between cosmic string cusp signals and blip glitches were considered from the point of view of the model. Because the model was not allowed to rely on coincidence, it was fully dependent on these morphologies, yielding unique insight. At the same time, this serves as a proof of concept for high classification performance before coincidence is introduced to further improve a pipeline.

There are several directions open for continued work, such as the inclusion of other gravitational-wave generating mechanisms on cosmic strings. These mechanisms are comprised by kinks and kink-kink collisions, signals with different spectral indices from cusps that are now starting to be considered in cosmic string searches [128,194]. Furthermore, there are indications the proposed deep-learning model offers additional room for improvement, for instance through an extended training phase. It is expected this adjustment will also serve to lower the false positive rate. In terms of analysis, the surgery process can be further refined, for instance by working at a resolution higher than three sections. This can be achieved by redefining the function of the standard deviation that marks the incisions.

Altogether, it is expected that the Einstein Telescope will bring a variety of new opportunities for the detection of cosmic strings and that deep learning will play a vital role in their analysis.

7.2 Intermediate-Mass Black Hole Binary Searches

Three different neural network classifiers of the same architecture were introduced, each with a different training regime that involves different levels of exposure to precession. It was shown that these classifiers perform exceptionally well for the task of telling signals from glitches, with the curriculum classifier in particular obtaining an accuracy as high as 95% on the held-out test sets. It was also shown that in a simulation on a test set of signals and glitches, all three classifiers outperform the matched filtering algorithm for the task of separating signals from glitches. Matched filtering was in particular hindered by a class of glitches with high values on the Wasserstein metric. The sensitivities of the classifiers to waveform parameters and statistics were investigated, finding that there are no particular combinations of these values that can accurately predict misclassification, although the three classifiers will become less confident as system total mass increases.

Through comparison of the three classifiers, the robustness of the neural network architecture to the effects of precession was tested. Although all three classifiers perform extremely well, it appears that the three classifiers handle precession differently and in accordance with expectations based on their provided training sets. This is measured through the effective precessing spin statistic, showing that the architecture is not fully robust to precession in the sense that the insertion of precession into the training phase will alter classifier behaviour. This is in part mitigated by the use of curriculum learning, which was shown to be able to improve classifier performance.

Future work could see the inclusion of additional waveform effects such as eccentricity or higher-order modes [34]. Alternatively, the architecture can be modified further. The early overfitting of the classifiers suggests that the architecture could benefit from the addition of max pooling layers to reduce the number of model parameters. There are other opportunities that can be found within the training regimes. One example can be the extension of curriculum learning to a curriculum that is strongly segmented over a larger number of phases, with examples of increasing difficulty being fed as the training progresses. This difficulty can be measured by for instance the effective precessing spin parameter, as was done in this work using a hard cut-off at zero.

This work has demonstrated that neural networks are efficient candidate models in the search for intermediate-mass black hole binaries in the presence of glitches. A baseline for the performance on this task has been established, while at the same time showing that neural networks may serve as a new medium to observe the sometimes minute differences between different classes of gravitational-wave signals and glitches. This exploration is important for future analyses, as such tasks have proven difficult for traditional methods. As these differences become better understood, the quality and quantity of gravitational-wave analyses should benefit directly, enabling a deeper understanding of the cosmos and the many objects residing in it.

7.3 Accelerating Waveform Generation

Evolutionary algorithms were introduced as a method for the optimisation problem of subsampling frequency spectra to speed-up waveform approximants, and per extension, the computation of SNR. It was shown that even with strict bounds on the percentual errors, speedups of at least 30% in computational time are achievable on nearly the entire sampled parameter ranges. If the error bound is heightened to at most 2.87% the frequency arrays generalise to the full parameter ranges (specified in Table 6.3). The margins of the parameter ranges where the frequency arrays were not effective were explored, as well as the possibility of tiling the parameter space to force a set upper bound on the error. It was demonstrated that the ranges where the frequency arrays were ineffective are mostly determined through correlation with the SNR, and it could therefore be argued that the frequency arrays generalise to the full ranges depending on the error metric in use. The proposed proof of concept frequency array, obtained by using GW150914 as a reference event, shows a speedup of 46% with a maximum error of 0.5% on the full parameter ranges. Additionally, it was shown that this frequency array performs exceptionally well for the reconstruction of the maxima of log likelihood functions.

Possible future lines of work include improving the methods presented in Ch. 6 specifically for parameter estimation [291] or time domain problems such as the fast computation of Euler angles that describe the evolution of precessing binary black hole systems [73,89,90,241]. Alternatively, the tiling of the parameter space, and simultaneously the optimum in trade-off between accuracy and generalisibility, can be studied further. The evolutionary algorithm can in theory also be calibrated to perform on targeted sections of the parameter space, such as those corresponding to intermediate-mass black hole binaries or extreme mass ratio inspirals. Within these sections a single frequency array might suffice.

Further improvements can be made to the evolutionary algorithm, for instance by incorporating more prior information. The algorithm can be discouraged from traversing sections of the solution space, for example by using a supplied power spectral density, or by more hyperparameter tuning. Currently the algorithm is hardcoded to run for a total of 30 generations. Despite a stopping condition being in place, the algorithm was ran for the full duration to study its behaviour. By tuning the hyperparameters and designing more advanced stopping conditions, the evolutionary algorithm can be modified to propose frequency arrays meeting a set of required conditions without additional user input. A future milestone is to release the adjustable code as a package.

For the evolutionary algorithm itself, the investigation of its applications to data compression for gravitational-wave data analysis could prove useful. This can be coupled to the exploration of alternative interpolation schemes.

With the number of gravitational-wave detections increasing through the arrival of detector upgrades and new generations of observatories, the need for faster yet accurate analyses is becoming more apparent. The use of optimisation schemes such as evolutionary algorithms to realise these targets are indispensable, and will enable the scientific community to learn more about the nature of gravitational waves in less time.

Acknowledgements

Cosmic String Searches

I thank Adrian Helmling-Cornell and the anonymous Physical Review D referee for their helpful comments. With thanks to Artim Bassant for being open to string-theoretical discussion. The authors of the corresponding paper Q.M and M.L. are supported by the research program of the Netherlands Organisation for Scientific Research (NWO). D.T. is supported by the Sherman Fairchild Postdoctoral Fellowship at Caltech. S.C. is supported by the National Science Foundation under Grant No. PHY-2309332. I am grateful for computational resources provided by the LIGO Laboratory and supported by the National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation.

Intermediate-Mass Black Hole Binary Searches

With thanks to Stefano Schmidt, Melissa Lopez, Bhooshan Gadre, Tom Dooney, Vasileios Skliris and Hugo Einsle. The authors of the corresponding paper Q.M and M.vd.S are supported by the research program of the Netherlands Organisation for Scientific Research (NWO). M.vd.S acknowledges support from the European Union, AHEAD 2020 (grant number 871158) and of the project Cortex with project number 00686766 of the research programme NWA which is (partly) financed by the Dutch Research Council (NWO). S.C is supported by the National Science Foundation under Grant No. PHY-2309332. I am grateful for computational resources provided by the LIGO Laboratory and supported by the National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation.

Accelerating Waveform Generation

With thanks to Harsh Narola, Stefano Schmidt, Melissa Lopez, Justin Janquart, Bhooshan Gadre, Marc van der Sluys, Chris Van Den Broeck, Justin Perez and the anonymous Physical Review D referee. I am supported by the research program of the Netherlands Organisation for Scientific Research (NWO). S.C is supported by the National Science Foundation under Grant No. PHY-2309332. I am grateful for computational resources provided by the LIGO Laboratory and supported by the National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation.

Public Summary

Einstein based his theories of gravity on, among others, the assumption that the speed of light is a *universal constant*, a measured physical constant that is the same to all who try to measure it. As a speed is expressed as a proportion of length to time, as for instance a car might travel at 60 km/h, there must exist a unique relation between time and the space where distance is measured. This relation is modelled in what is called a *spacetime*. As the use of a compound word suggests, the spacetime does exactly what one would expect it to do – it is used to model not only space, but simultaneously models time. This requires the use of four dimensions, with three dimensions reserved for locations measured in space, and one for locations measured in time.

Although the concept of a dimension may sound complicated, it does not have to be. A *dimension* is in essence no more than an axis for measurement. In any space imaginable, such as a living room with a door for reference, the spatial location of an object can be specified by three measurements. This corresponds to three dimensions. Starting measurements from the doorframe, a hanging wall clock might be one meter forwards, two meters to the right, and at a height of two meters. This gives three axes of measurement, or a measurement in three dimensions. The dimension of time is slightly more abstract. Since we perceive ourselves as moving in three dimensions, as we would in the living room, it is difficult to picture the fourth dimension. Its measurement can be pictured however, simply by reading the time measurement as given by the wall clock. Combining this with the space measurement gives four measurements that uniquely specify a location in both time and space, or in spacetime. The specification of such locations does not require a physical clock to be present, and this in fact happens on a daily basis, for instance when an appointment is written down in an agenda. The appointment will be at a given spatial address, and at a given time. Therefore, an agenda can be thought of as registering spacetime locations.

Spacetime as an abstract model tells us little about why its design would be a requirement for the formulation of a theory of gravity, or how gravity would even act within spacetime. This is because gravity is understood through the matter that is present within spacetime. A *black hole*, an astrophysical object so extreme that it can imprint noticeable effects on spacetime, can be used to illustrate this relationship. If different clocks were to be scattered at different distances around a black hole, each of these clocks would measure the passing of time differently. This is because of a phenomenon called *time dilation*, causing the time component of a spacetime measurement to become relative and depending on the location of measurement. The extreme properties of the black hole, in particular the extreme density of mass, have deformed the region of spacetime surrounding the black hole. This deformation is not only measurable in the time dimension, but also in the dimensions of space. Objects traveling in the vicinity of the black hole will experience a gravitational pull from the black hole, which can alter the paths of the objects. An example of such an effect is noticeable in every day life, as it is the comparatively miniscule attraction of Earth's gravity that keeps us tethered to its surface. According to *general relativity*, named as such because all measurements are evidently relative, such deformations of spacetime are what we know as gravity. In the example of a black hole, the mass of the black hole has dictated the presence of gravity, and gravity will dictate the movement of passing objects. It is important to note that these objects might have mass too, and might therefore have their own effects on spacetime. This illustrates the complicated relationships between gravity and matter that is innate to general relativity.

The number of dimensions in a model can be reduced by ignoring measurements. Looking straight down, the height of an object might not be important, and under this assumption two dimensions for space measurements will suffice. A similarly lower-dimensional aid for the visualisation of spacetime deformation could be a bowling ball resting on the mat of a trampoline, sinking down in its surface. Analogously, a black hole can be pictured as being suspended in spacetime, weighing down the fabric of spacetime itself. Although for the example of the bowling ball time dilation will not be noticeable, the presence of the heavy object will certainly impact the trajectory of marbles rolling over the mat. If the bowling ball is moved or lifted, the imprint on the mat will change, as the tension of the mat will necessarily cause it to follow the mass. This shows that it is not enough to understand static masses in order to fully understand gravity, and that in fact moving and accelerating masses need to be considered as well. In the latter case, the acceleration of mass generates waves. These waves will behave similarly as waves in water, but propagated in the medium of spacetime. Such waves are called *gravitational waves*. Gravity, or the deformations of spacetime, can change periodically as gravitational waves pass. Similar changes can be intuitively recognised in roller coasters, where a rider might be pushed down into their seat at one point, and nearly be lifted out from weightlessness at another. So too gravity might change at the location where a gravitational wave is measured.

Important processes that generate gravitational waves measurable on Earth are mergers of *compact objects*, objects of very high mass density of which black holes and *neutron stars* are examples. Neutron stars are stars with properties almost as extreme as those of black holes, and consisting at least in large part of neutrons. During a merger two such objects in a *binary* will spiral closer together on accelerating trajectories before forcibly merging. These extreme processes are extremely far removed from Earth, and since the intensity of waves must decrease as they travel, we do not consciously experience the passing of these gravitational waves.

In fact, by the time the gravitational waves reach Earth, they become extremely difficult to measure at all, even in dedicated gravitational-wave observatories. It is, however, possible. Laser-based detectors use *interferometers* that search for passing gravitational waves by measuring extremely small differences in laser arrival times. Broadly speaking, if the arrival times of the electromagnetic waves emitted by the very precise lasers change, the distance traveled by these electromagnetic waves must have changed. This betrays a deformation of spacetime at the position of the observatory. By registering such patterns at scales as small of a fraction of a proton, gravitational waves can be identified and analysed. By comparing the data to *waveforms*, known scientific models of what gravitational waves should look like, scientists can discern the specific details of the events that generated the gravitational waves. This is a part of *gravitational-wave astronomy.*

As gravitational waves are minute, these observatories are necessarily extremely sensitive. This has the unwanted downside that the observatories measure a large amount of data that does not actually contain any gravitational-wave signals. This is referred to as detector noise, or *noise* for short, defined as patterns in the data that are not generated by gravitational waves. Returning noise patterns are called detector glitches, or simply *glitches.* Glitches are undesirable in particular, as they may show recurring patterns that are very reminiscent of gravitational-wave signals. In order to tell glitches from signals, scientists employ and have developed advanced data analysis methods based on statistics and machine learning. Here, *machine learning* is understood as the subfield of data analysis where machines are made to independently attempt the learning of patterns.

The research presented in this thesis is centered around two problems in *gravitational-wave searches.* The first of these is the answering of the question of how machine learning methods can be employed to accurately tell signals from glitches, and how the machine learning models that are becoming exceedingly good at this task can teach the scientific community more about the differences between the two. Signals and glitches can not typically be told apart by eye, and in many cases, traditional data analysis methods fail as well. If machine learning models can recognise discriminating patterns, it would be wise to learn what it is that the machine has learned.

In pursuit of addressing the first problem this methodology was applied to *cosmic string* searches. Cosmic strings are cosmological objects that have not yet been detected but that do possibly exist, supported by the existence of very similar objects existing in other theories, with the theory of electromagnetism as an example. As it is known what cosmic string signals should look like, the detector data can be searched for these signals. Unfortunately, it is also known that these signals will look almost exactly like a very prolific class of glitches. Since finding a cosmic string signal would be a first, it is important that no signals are discarded because they are mistaken for glitches, as the discovery of cosmic strings will reveal new physics. The result of this research is a *neural network*, a machine learning model inspired by a biological model of the brain, that was shown to be very efficient at telling cosmic string signals from glitches. It was also shown in specifically designed tests to outperform the current state-of-the-art. Investigating this complicated model, it was found that there are certain properties that may help tell the difference between the cosmic string signals and glitches without the mediation of machine learning. In machine learning this is referred to as *interpretability*, the study of cause and effect in a classification model. In the case at hand, the causes are slight differences between the signals and glitches, and the effects are different classifications.

Similar neural networks were trained for the problem of telling the signals from *intermediate-mass black hole binaries* from glitches. Black holes can be classified based on their mass as measured in *solar masses*, expressing the mass of the black holes in terms of the mass of the sun. The category of intermediate-mass black holes, so called because of its positioning between the categories of stellar-mass black holes and supermassive black holes, is considered in particular. The reason for this is that their waveforms will show great resemblance to glitches. At the time of writing, only one such system has been observed through gravitational waves, meaning that here too it is of the essence to ensure signals are not confused for glitches. Making such mistakes could lead to misconceptions about the population of intermediate-mass black holes, which may evolve into supermassive black holes. The latter are known to drive galaxy formation and evolution, making intermediate-mass black holes important per extension. The obtained neural networks were found to perform extremely well. The interpretation of these models is however different to the models obtained for cosmic string searches, as a black hole binary can exhibit behaviour that a singular cosmic string can not. The particular behaviour considered in this research is that of *precession*, a dynamic change of the black hole trajectories that affects the waveform. Modelling this effect requires several parameters. It was found through extensive analysis what impact the inclusion of precession has on the neural networks, and that the neural networks performed well in test cases involving precession even if the effects of precession were not explicitly taught. However, teaching precession to the model following the right curriculum can still boost performance. Like for cosmic strings, the neural networks again outperformed the state-of-the-art in the tested scenarios. In spite of this it was derived that the neural networks do seem to show similar characteristics and sensitivities as the traditional methods do.

The second problem is the fast generation of waveforms. In gravitational-wave applications, models for waveforms are typically computed part-wise, computing the power at every individual *frequency* of the signal. This is possible because like music and more generally sound waves, gravitational-wave signals can be decomposed into contents at different frequencies. Since a large amount of waveforms is typically generated for analyses, and because this generation is usually one of the most costly steps, reducing the number of frequency contents that need to be computed can significantly decrease time spent. The research assumes that not all frequency contents need to be computed in the first place, and instead, at a selection of frequencies the required values can be determined from the nearby frequencies. This would be a lot faster than directly computing the actual contents. It would then need to be found for what frequencies these assumptions do indeed hold, as interpolating at the wrong frequencies may significantly impact the accuracy of the computed waveform. This question was answered with the help of *evolutionary algorithms*, algorithms that attempt to solve difficult problems using Darwin's principle of natural selection. As happens in nature according to this principle, solutions are continuously tested and improved by letting the best solutions parent new solutions in the hope that they outperform their parents. Through the use of such algorithms it was shown that it is possible to compute waveforms significantly faster at only a very minor loss in accuracy. It was also shown that the developed methods are applicable for the particular application of *parameter estimation*, where system parameters are estimated from gravitational-wave data after a signal has been identified.

The found answers to these questions were released in three papers presented as three separate chapters in this thesis, along with the necessary backgrounds in phyics, statistics, and machine learning.

Openbare Samenvatting

De theorie van zwaartekracht zoals deze geformuleerd is door Einstein berust onder andere op de aanname dat de snelheid van het licht een *universele constante* is. Dit betekent dat deze snelheid fysiek meetbaar is en altijd dezelfde waarde aan moet nemen, ongeacht waar, wanneer, of door wie deze gemeten wordt. Gezien snelheden worden uitgedrukt als proporties van afstand tot tijd, zoals bijvoorbeeld een auto die 60 km/u rijdt, betekent dit dat er een bijzondere relatie moet bestaan tussen tijd en de ruimte waarin de afstand gemeten wordt. Deze relatie wordt gemodelleerd in een *ruimtetijd*. De naam van dit model wordt gegeven door een samengesteld woord, en een ruimtetijd doet precies wat deze samenstelling doet vermoeden. De ruimtetijd wordt niet alleen gebruikt als model voor de ruimte, maar modelleert tegelijkertijd ook het verloop van tijd in de ruimte. Hiervoor zijn in totaal vier dimensies vereist. Drie van deze dimensies zijn gereserveerd voor ruimtelijke locaties, en de vierde dimensie wordt gebruikt voor de meting van tijd.

Hoewel de term dimensie lastig kan klinken hoeft het idee niet lastig te begrijpen te zijn. Een *dimensie* is niet meer dan een as waarlangs gemeten wordt. Een voorbeeld kan gegeven worden door een woonkamer, waarin de deur gebruikt wordt als ijkpunt. In deze ruimte kan de locatie van een object gegeven worden door drie metingen. Beginnend bij de deurpost kan de locatie van een wandklok beschreven worden als zijnde een meter voorwaards, twee meter naar rechts, en op een hoogte van twee meter. Dit geeft een meting langs drie assen, ofwel een meting in drie dimensies. Gezien mensen zich bewust bewegen in slechts drie fysieke dimensies is het lastiger een inbeelding te maken van een vierde dimensie voor tijd. Een tijdsmeting kan echter gemakkelijk aan het vorige voorbeeld toegevoegd worden door de tijd op de klok af te lezen. Gecombineerd met de locatie van de klok geeft dit vier metingen die uniek vast leggen waar in de ruimte en tijd, ofwel in de ruimtetijd, de klok zich bevindt. Het is belangrijk om in te zien dat de aanwezigheid van een klok geen vereiste is om de tijd aan te geven, en dat dit geregeld gebeurt in het dagelijks leven. Wanneer een afspraak in een agenda genoteerd wordt is dit vaak in de vorm van een adres en een tijd. Een alledaagse agenda kan daarom ook gezien worden als een lijst van ruimtetijd locaties.

Op dit punt vertelt een ruimtetijd als abstract model nog niets over waarom de ruimtetijd nodig is om zwaartekracht te kunnen begrijpen, of hoe zwaartekracht zich überhaupt manifesteert in dit model. Dit is omdat zwaartekracht begrepen wordt door middel van de materie die aanwezig is in de ruimtetijd. Een *zwart gat* is een astrofysisch object met zulke extreme eigenschappen dat het zwarte gat de ruimtetijd meetbaar aan kan tasten. Door deze extreme effecten vormen zwarte gaten een goed voorbeeld aan de hand waarvan deze relatie uitgelegd kan worden. Als een zwart gat overwogen wordt, en er rondom het zwarte gat op verschillende afstanden klokken geplaatst zouden worden, zou iedere klok het verloop van tijd anders meten. Dit is het gevolg van een verschijnsel genaamd *tijddilatie*, wat betekent dat de meting van tijd relatief wordt tot de locatie van de meting. In dit geval heeft de extreme dichtheid van de materie van het zwarte gat de nabije ruimtetijd dusdanig aangetast dat de tijdsdimensie vervormd is geraakt. Dit soort effecten beperken zich niet alleen tot tijd, maar tonen zichzelf ook in de ruimtelijke dimensies. Objecten die zich langs het zwarte gat bewegen zullen een aantrekkingskracht van het zwarte gat ervaren die hun trajecten aan kunnen passen. Ook hiervan zijn voorbeelden zichtbaar in ons dagelijks leven, zoals dat het de relatief lichte zwaartekracht van de aarde is die ervoor zorgt dat we veilig op het oppervlak van de aarde blijven. Volgens *algemene relativiteit*, zo genoemd omdat de metingen derhalve relatief zijn, zijn dit soort vervormingen een en hetzelfde als zwaartekracht. In het voorbeeld van het zwarte gat is het de massa die de aanwezigheid van zwaartekracht dicteert, en dicteert de zwaartekracht de bewegingen van nabije objecten. Het is belangrijk om er bij stil te staan dat deze objecten zelf ook massa kunnen bezitten, en daarom hun eigen effecten uit kunnen oefenen op de ruimtetijd. Dit laat zien hoe lastig de relaties tussen zwaartekracht en materie kunnen zijn binnen de theorie van algemene relativiteit.

Door metingen te negeren kan het aantal te overwegen dimensies worden verminderd. Wanneer van te voren bekend is dat een meting altijd gedaan zal worden van recht boven een object zal de hoogte van het object niet uit maken, gezien deze toch niet zichtbaar zal zijn. Dit betekent dat onder deze aanname twee dimensies voor de ruimtelijke locatie zullen volstaan. Een soortgelijke lager-dimensionale hulp voor de inbeelding van ruimtetijd vervormingen kan gegeven worden door een bowling bal die weg gezakt is in de mat van een trampoline en daar stil ligt. Analoog kan een zwart gat gezien worden als zijnde hangend in de ruimtetijd, waar deze het weefsel van de ruimtetijd als het ware omlaag trekt. In het geval van de bowling bal zal tijddilatie niet merkbaar zijn, maar de aanwezigheid op de mat zal zeker impact hebben op de trajecten van knikkers die over de mat gerold kunnen worden. Als de bowling bal nu echter opgelicht wordt zal de mat aantrekken doordat de spanning er voor zorgt dat de mat de massa van de bal moet volgen. Dit laat zien dat het niet genoeg is om statische massas te begrijpen om een volledige grip te krijgen op zwaartekracht, maar dat ook bewegende en versnellende massas overwogen moeten worden. In het laatste geval zal de versnelling van de massa golven genereren op dezelfde manier als golven in water gegenereerd worden, maar dan in de ruimtetijd. Deze golven worden *zwaartekrachtgolven* genoemd. Zwaartekracht, of de vervorming van ruimtetijd, kan veranderen als een zwaartekrachtgolf passeert. Soortgelijke veranderingen kunnen intuitief herkend worden in achtbanen, waar de inzittenden op een punt in hun stoelen gedrukt kunnen worden, en op een ander moment alleen gestopt worden door hun beugels vanwege tijdelijke gewichtloosheid. Op dezelfde manier kan de zwaartekracht veranderen op de plek waar een zwaartekrachtgolf wordt gemeten.

Belangrijke processen die op aarde meetbare zwaartekrachtgolven produceren zijn de versmeltingen van *compacte objecten*, objecten met zeer hoge massa dichtheid waar zwarte gaten en *neutronen sterren* specifieke voorbeelden van zijn. Neutronensterren zijn sterren met eigenschappen die bijna zo extreem zijn als de eigenschappen van zwarte gaten, en die op zijn minst grotendeels bestaan uit neutronen. Een versmelting vindt plaats wanneer twee zulke objecten versnellend om elkaar heen bewegen in een *binair systeem*, constant dichter naar elkaar toe trekkend tot ze geforceerd worden samen te komen. Dit soort extreme processen vinden plaats op extreem verre afstand van de aarde, en gezien de intensiteit van golven afnemen als ze afstand afleggen worden deze zwaartekrachtgolven niet door mensen opgemerkt wanneer ze ons passeren. De intensiteit neemt zelfs zo ver af dat het extreem moeilijk is om de zwaartekrachtgolven met inspanning te kunnen meten, zelfs met gespecialiseerde instrumenten. Moeilijk betekent echter niet onmogelijk. Een *interferometer* is een instrument dat lasers gebruikt om passerende zwaartekrachtgolven te meten door zeer kleine afwijkingen in aankomsttijden van de laserstralen te meten. Het gedachtenproces is dat wanneer de aankomsttijden van de electromagnetische straling uitgezonden door de erg precieze lasers verandert, de afstand afgelegd door deze straling veranderd moet zijn. Dit betekent dat de ruimtetijd vervormd moet zijn op de locatie van het observatorium. Door zulke patronen, gemeten op schalen zo klein als een fractie van een proton, te registreren kunnen zwaartekrachtgolven gemeten en geanalyseerd worden. Door deze data te vergelijken met *golfvormen*, wetenschappelijke modellen van hoe de wetenschap verwacht dat zwaartekrachtgolven van specifieke afkomst er uit zouden moeten zien, is het mogelijk de details van het evenement dat de zwaartekrachtgolven heeft gegenereerd aan de data te ontfutselen. Dit maakt deel uit van *zwaartekrachtgolven astronomie*.

Gezien de miniscule intensiteit van de zwaartekrachtgolven die waarneembaar zijn op aarde moeten de observatoria extreem gevoelig zijn. Dit heeft als ongewenst bijeffect dat een hoop data vastgelegd wordt waar geen werkelijke signalen van zwaartekrachtgolven in zitten. Deze waarnemingen worden *ruis* genoemd, gedefinieerd als patronen in de data die niet gegenereerd zijn door zwaartekrachtgolven. Herhalende patronen worden *artefacten* genoemd. Artefacten zijn specifiek ongewenst omdat ze erg veel kunnen lijken op de signalen van werkelijke zwaartekrachtgolven. Om deze artefacten van signalen te kunnen onderscheiden gebruikt de gemeenschap geavanceerde data analyse methoden gebaseerd op statistiek en machine learning. In dit geval wordt met *machine learning* het gebied aangeduid waar machines onafhankelijk patronen leren te herkennen, als deelgebied van data analyse.

Het onderzoek dat gepresenteerd wordt in dit proefschrift richt zich op twee problemen in de *opsporing van zwaartekrachtgolven*. Het eerste probleem wordt gegeven door de vraag hoe machine learning toegepast kan worden om signalen van artefacten te onderscheiden, en hoe de continu verbeterend machine learning modellen gebruikt kunnen worden om de gemeenschap meer te leren over de verschillen tussen deze twee klassen. Typisch kunnen signalen en artefacten niet onderscheiden worden met het blote oog, en in veel gevallen slagen ook traditionele data analyse methoden hier niet in. Als machine learning modellen dit verschil wel consistent kunnen leren doet de gemeenschap er verstandig aan lering te trekken uit wat het is dat de machines leren.

Om dit eerste probleem te adresseren is deze methodologie toegepast in de opsporing van *kosmische snaren*. Dit zijn kosmologische objecten die nog niet waargenomen zijn, maar die mogelijk bestaan. Dit mogelijke bestaan wordt onderbouwd door het bestaan van soortgelijke verschijnselen in andere theorien, bijvoorbeeld in de theorie van electromagnetisme. Gezien het bekend is hoe de golfvormen van kosmische snaren er uit zouden moeten zien kan er in de data gezocht worden naar signalen van kosmische snaren. Echter wil het lot dat deze signalen zeer veel lijken op een zeer veelvoorkomende klasse artefacten. Gezien kosmische snaren nog niet eerder waargenomen zijn is het belangrijk dat signalen niet ten onrechte genegeerd worden omdat ze verward worden voor een artefact, gezien de observatie zou leiden tot verdere ontwikkeling van de natuurkunde. Het resultaat van dit onderzoek is een *neuraal netwerk*, een machine learning model dat is geinspireerd door een biologisch model van de hersenen. In de loop van het onderzoek is laten zien dat dit neurale netwerk zeer effectief is in het onderscheiden van kosmische snaar signalen van artefacten, en dat het in een gecontroleerde omgeving zelfs beter presteert dan de huidige standaard in het veld. Door dit model te onderzoeken is gevonden dat er specifieke eigenschappen in de golfvormen zijn die gebruikt kunnen worden om het verschil tussen de signalen en artefacten te zien zonder de tussenkomst van machine learning. In machine learning wordt dit soort onderzoek aangeduid als *interpreteerbaarheid*, wat de studie van oorzaak en gevolg is in de context van een classificatie model. In het huidige geval zijn de oorzaken kleine verschillen tussen signalen en artefacten, en de effecten verschillende classificaties.

Soortgelijke neurale netwerken worden in dit proefschrift onderzocht om signalen van *binaire middelzware zwarte gaten* van artefacten te kunnen onderscheiden. Zwarte gaten kunnen in categoriën ingedeeld worden op basis van hun massa. Deze massa wordt uitgedrukt in zonnemassas, ofwel in termen van de massa van de zon. De categorie van middelzware zwarte gaten wordt dan zo genoemd omdat deze categorie gepositioneerd is tussen de categoriën van de stellaire zwarte gaten en superzware zwarte gaten. De categorie van binaire middelzware zwarte gaten wordt specifiek overwogen omdat de golfvormen van deze systemen zeer veel lijken op die van artefacten. Op het moment van schrijven is er maar een enkel binair middelzwaar zwart gat waargenomen, wat betekent dat het ook hier belangrijk is dat de signalen niet verward worden voor artefacten, gezien dit zou kunnen leiden tot misvattingen over de spreiding van middelzware zwarte gaten. Deze spreiding is belangrijk om te begrijpen omdat de middelzware zwarte gaten uit kunnen groeien tot superzware gaten, die op hun beurt de ontwikkeling en groei van sterrenstelsels drijven. Dit maakt het onderzoek van middelzware zwarte gaten zeer belangrijk. De neurale netwerken die zijn opgeleverd als resultaat van dit onderzoek presteren extreem goed. De interpretatie van deze neurale netwerken is echter anders dan de netwerken die voor kosmische snaren ontwikkeld zijn, gezien binaire zwarte gaten verschijnselen kunnen laten zien die niet gezien worden bij enkelvoudige kosmische snaren. Het specifieke verschijnsel dat overwogen wordt is *precessie*, dynamische afwijkingen van de paden van de zwarte gaten die de golfvorm aantasten. Om dit effect te kunnen modelleren zijn meerdere parameters nodig, wat de golfvorm modellen complexer maakt dan die van kosmische snaren. Door uitbundige analyse is gevonden welke impact de toevoeging van precessie aan de golfvormen heeft op de neurale netwerken, en dat de neurale netwerken zeer goed presteren zelfs wanneer de effecten van precessie niet expliciet vooraf geleerd zijn. Het is verder aangetoond dat het vooraf leren van precessie de prestaties nog steeds kan verbeteren. Zoals in het geval van de kosmische snaren is laten zien dat ook deze neurale netwerken in een gecontroleerde omgeving beter presteren dan de huidige standaard. Het is echter ook gedemonstreerd dat de neurale netwerken dezelfde karakteristieken als deze standaard gebruiken om onderscheid te maken tussen signalen en artefacten, en daarom gevoelig zijn voor verschillen in dezelfde parameters.

Het tweede probleem geadresseerd door het onderzoek in dit proefschrift is de versnelling van het aanmaken van golfvormen. In veel toepassingen voor zwaartekrachtgolven worden deze golfvorm modellen deelsgewijs berekend door de inhoud per *frequentie* van het signaal te berekenen. Net zoals voor muziek, en voor geluidsgolven in het algemeen, is het mogelijk om zwaartekrachtgolven te ontbinden in verschillende frequenties. Gezien voor typische analyses grote hoeveelheden golfvormen gegenereerd moeten worden, en omdat dit typisch een van de meest kostbare stappen van het proces is, kan het terugbrengen van het aantal frequenties dat overwogen moet worden er voor zorgen dat de benodigde tijd sterk verminderd wordt. Het onderzoek gaat er van uit dat de inhoud van niet alle frequenties expliciet berekend hoeft te worden, maar dat in plaats daarvan de benodigde waarden bepaald kunnen worden uit soortgelijke frequenties. Dit is een stuk sneller dan de nodige waarden op iedere frequentie berekenen. Wanneer dit gedaan wordt is het nodig te bepalen op welke frequenties deze aanname geldig is, gezien interpoleren op de verkeerde frequenties een significante impact op de accuraatheid van de golfvorm als gevolg kan hebben. Waar dit gedaan kan worden is gedemonstreerd door de toepassing van *evolutionaire algoritmen*. Dit zijn algoritmen die moeilijke problemen oplossen door gebruik te maken van het principe van natuurlijke selectie, zoals voorgesteld door Darwin. Dit principe stelt dat organismen oplossingen zijn voor problemen van hun leefomgevingen, en daar constant getest en verbeterd worden door alleen de beste oplossingen nageslacht te laten produceren. De verwachting is dat het nageslacht verder zal verbeteren op de oplossing zoals deze geboden is door de ouders. Door toepassing van dit type algoritmen is laten zien dat het mogelijk is golfvormen significant sneller te berekenen, tegen een klein offer in accuraatheid. Het is verder laten zien dat de ontwikkelde methoden van toepassing zijn op het deelgebied van *parameterschatting*, waar systeem parameters geschat worden op basis van de data nadat een signaal in de data is waargenomen.

De gevonden antwoorden op deze problemen zijn uitgebracht in drie papers die in dit proefschrift corresponderen met drie verschillende hoofdstukken. Deze hoofdstukken worden ingeleid met de nodige achtergronden in de natuurkunde, statistiek en machine learning.

Curriculum Vitae

Personal Information

Name	Reinier Henricus Antonius Johannes (Quirijn) Meijer
Nationality	Dutch
Date of Birth	September 1st, 1992
Email	r.h.a.j.meijer@uu.nl

Education

2020 - 2024	PhD in Physics, *Utrecht University*
2017 - 2019	MSc in Computer Science, *Utrecht University*
2013 - 2017	BSc in Mathematics, *Utrecht University*

Publications

- Q. Meijer, M. Lopez, D. Tsuna and S. Caudill. "Gravitational-Wave Searches for Cosmic String Cusps in Einstein Telescope Data using Deep Learning", *Physical Review D*, 109(2): 22006-22021, January 2024.
- Q. Meijer and S. Caudill. "Fast Waveform Generation for Gravitational Waves using Evolutionary Algorithms", *Physical Review D*, 110(4): 43035-43050, August 2024.
- A. Martins, M. Lopez, Q. Meijer, G. Balthus, M. van der Sluys, C. Van Den Broeck and S. Caudill. "Improving Early Detection of Gravitational Waves from Binary Neutron Stars using CNNs and FPGAs", *arXiv e-Print Archive*, arXiv:2409.05068 [astro-ph.IM], September 2024.
- Q. Meijer, M. van der Sluys and S. Caudill. "Robustness of Deep Learning Models to Precession in Gravitational-Wave Searches for Intermediate-Mass Black Hole Binaries", *arXiv e-Print Archive*, arXiv:2410.07855 [gr-qc], October 2024.
- H. Narola, J. Janquart, Q. Meijer, K. Haris and C. Van Den Broeck. "Gravitational-Wave Parameter Estimation with Relative Binning: Inclusion of Higher-Order Modes and Precession, and Applications to Lensing and Third-Generation Detectors", *Physical Review D*, 110(8): 84085-84100, October 2024.

Bibliography

[1] J. P. Luminet. Image of a spherical black hole with thin accretion disk. *Astronomy and Physics*, 75:228–235, May 1979.

[2] C. E. Kohlhase and P. A. Penzo. Voyager mission description. *Space Science Reviews*, 21(2):77–101, Nov 1977.

[3] Neil Ashby. Relativity in the Global Positioning System. *Living Rev. Rel.*, 6:1, 2003.

[4] Isaac Newton. *Philosophiae Naturalis Principia Mathematica.* 1687.

[5] A. Einstein. Zur Elektrodynamik bewegter Körper. *Annalen der Physik*, 322(10):891–921, January 1905.

[6] H. Minkowski. Die grundgleichungen für die elektromagnetischen vorgänge in bewegten körpern. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1908:53–111, 1908.

[7] Albert Einstein. Die Feldgleichungen der Gravitation. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften*, pages 844–847, January 1915.

[8] F. W. Dyson, A. S. Eddington, and C. Davidson. A Determination of the Deflection of Light by the Sun's Gravitational Field, from Observations Made at the Total Eclipse of May 29, 1919. *Philosophical Transactions of the Royal Society of London Series A*, 220:291–333, January 1920.

[9] Barbara Ryden. *Introduction to Cosmology.* 2016.

[10] S.M. Carroll. *Spacetime and Geometry.* Cambridge University Press, 2019.

[11] Oliver Heaviside. A Gravitational and Electromagnetic Analogy. *The Electrician*, 31:281–282, 1893.

[12] Henri Poincaré. Sur la Dynamique de l'Electron. *Comptes Rendus*, 140:1504–1508, 1905.

[13] Albert Einstein. Über Gravitationswellen. *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)*, 1918:154–167, 1918.

[14] J. H. Taylor, L. A. Fowler, and P. M. McCulloch. Measurements of general relativistic effects in the binary pulsar PSR 1913+16. *Nature*, 277:437–440, 1979.

[15] J. M. Weisberg, J. H. Taylor, and L. A. Fowler. Gravitational Waves from an Orbiting Pulsar. *Sci. Am.*, 245:66–74, 1981.

[16] Hermann Bondi. Plane gravitational waves in general relativity. *Nature*, 179:1072–1073, 1957.

[17] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, B. Allen, A. Allocca, P. A. Altin, S. B. Anderson, W. G. Anderson, K. Arai, and M. A. Arain. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.

[18] K. Thorne and C. Nolan. *The Science of Interstellar*. W. W. Norton, 2014.

[19] Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, Lindy Blackburn, Wilfred Boland, Katherine L. Bouman, and Geoffrey C. Bower. First m87 event horizon telescope results. i. the shadow of the supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L1, April 2019.

[20] Lia Medeiros, Dimitrios Psaltis, Tod R. Lauer, and Feryal Özel. The image of the m87 black hole reconstructed with primo. *The Astrophysical Journal Letters*, 947(1):L7, April 2023.

[21] Jerome Kristian, Robert Joseph Brucato, Natarajan Visvanathan, Howard H. Lanning, and Allan R. Sandage. On the optical identification of cygnus x-1. *The Astrophysical Journal*, 168, 1971.

[22] T. W. B. Kibble. Topology of Cosmic Domains and Strings. *J. Phys. A*, 9:1387–1398, 1976.

[23] Mark Wyman, Levon Pogosian, and Ira Wasserman. Bounds on cosmic strings from wmap and sdss. *Phys. Rev. D*, 72:023513, Jul 2005.

[24] Kun Qin, Sanyou Zeng, Zhengjun Li, and Hongyong Jing. Automated antenna design using self-adaptive differential evolution algorithm. In *2011 Fourth International Conference on Intelligent Computation Technology and Automation*, volume 2, pages 964–969, 2011.

[25] Quirijn Meijer, Melissa Lopez, Daichi Tsuna, and Sarah Caudill. Gravitational-wave searches for cosmic string cusps in einstein telescope data using deep learning. *Phys. Rev. D*, 109:022006, Jan 2024.

[26] Quirijn Meijer, Marc van der Sluys, and Sarah Caudill. Robustness of deep learning models to precession in gravitational-wave searches for intermediate-mass black hole binaries, Oct 2024. *arXiv e-Print Archive*, arXiv:2410.07855 [gr-qc].

[27] Quirijn Meijer and Sarah Caudill. Fast waveform generation for gravitational waves using evolutionary algorithms. *Phys. Rev. D*, 110:043035, Aug 2024.

[28] G.L. Naber. *The Geometry of Minkowski Spacetime: An Introduction to the Mathematics of the Special Theory of Relativity.* Applied mathematics sciences. Dover Publications, 2003.

[29] J.M. Lee. *Introduction to Riemannian Manifolds.* Graduate Texts in Mathematics. Springer International Publishing, 2019.

[30] L. Godinho and J. Natário. *An Introduction to Riemannian Geometry: With Applications to Mechanics and Relativity.* Universitext. Springer International Publishing, 2014.

[31] R.M. Wald. *General Relativity.* University of Chicago Press, 2010.

[32] S. Hawking and G.F.R. Ellis. *The Large Scale Structure of Space-Time.* Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1973.

[33] R.K. Sachs and H.H. Wu. *General Relativity for Mathematicians.* Graduate Texts in Mathematics. Springer New York, 2012.

[34] M. Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments.* Gravitational Waves. OUP Oxford, 2008.

[35] J. Lee. *Introduction to Smooth Manifolds.* Graduate Texts in Mathematics. Springer New York, 2012.

[36] C.W. Misner, K.S. Thorne, J.A. Wheeler, and D.I. Kaiser. *Gravitation.* Princeton University Press, 2017.

[37] I.G. Petrovsky. *Lectures on Partial Differential Equations.* Dover Books on Mathematics. Dover Publications, 2012.

[38] L.C. Evans. *Partial Differential Equations.* Graduate studies in mathematics. American Mathematical Society, 2010.

[39] Arthur M. Jaffe. The Millennium Grand Challenge in Mathematics. *Not. Amer. Math. Soc.*, 53(6):652–660, 2006.

[40] Keith Riles. Recent searches for continuous gravitational waves. *Modern Physics Letters A*, 32(39):1730035, December 2017.

[41] B. P. Abbott et al. All-sky search for short gravitational-wave bursts in the second advanced ligo and advanced virgo run. *Phys. Rev. D*, 100:024017, Jul 2019.

[42] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, S. Akcay, T. Akutsu, S. Albanesi, A. Allocca, P. A. Altin, A. Amato, C. Anand, S. Anand, A. Ananyeva, S. B. Anderson, W. G. Anderson, and M. Ando. 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.

[43] Nelson Christensen. Stochastic gravitational wave backgrounds. *Reports on Progress in Physics*, 82(1):016903, November 2018.

[44] Albert Abraham Michelson and Edward W. Morley. On the relative motion of the earth and the luminiferous ether. *American Journal of Science*, s3-34:333 – 345, 1887.

[45] A. Perot and Charles Fabry. On the application of interference phenomena to the solution of various problems of spectroscopy and metrology. *The Astrophysical Journal*, 9:87, 1899.

[46] B P Abbott, R Abbott, T D Abbott, S Abraham, F Acernese, K Ackley, C Adams, V B Adya, C Affeldt, M Agathos, K Agatsuma, N Aggarwal, O D Aguiar, L Aiello, A Ain, P Ajith, T Alford, G Allen, A Allocca, M A Aloy, P A Altin, A Amato, A Ananyeva, S B Anderson, W G Anderson, S V Angelova, S Antier, S Appert, K Arai, M C Araya, J S Areeda, M Arène, and N Arnaud. A guide to ligo–virgo detector noise and extraction of transient gravitational-wave signals. *Classical and Quantum Gravity*, 37(5):055002, February 2020.

[47] Interview with Nikhef's Mechanical Designer Marco Kraan: "I See Things that are too Beautiful Not to Share". `https://www.nikhef.nl/en/focusblokken/interview-with-nikhefs-mechanical-designer-marco-kraan-i-see-things-that-are-` Accessed on 04-10-2024.

[48] Gravitational Waves from a Binary Black Hole Merger Observed by LIGO and Virgo. `https://news.mit.edu/2017/gravitational-waves-binary-black-hole-merger-observed-ligo-and-virgo-0927`. Accessed on 04-10-2024.

[49] J. Aasi et al. Advanced ligo. *Classical and Quantum Gravity*, 32(7):074001, mar 2015.

[50] F. Acernese et al. Advanced virgo: A second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, dec 2014.

[51] K L Dooley, J R Leong, T Adams, C Affeldt, A Bisht, C Bogan, J Degallaix, C Graf, S Hild, J Hough, A Khalaidovski, N Lastzka, J Lough, H Luck, D Macleod, L Nuttall, M Prijatelj, R Schnabel, E Schreiber, J Slutsky, B Sorazu, K A Strain, H Vahlbruch, M Was, B Willke, H Wittel, K Danzmann, and H Grote. Geo 600 and the geo-hf upgrade program: successes and challenges. *Classical and Quantum Gravity*, 33(7):075009, March 2016.

[52] T Akutsu, M Ando, K Arai, Y Arai, S Araki, A Araya, N Aritomi, Y Aso, S Bae, Y Bae, L Baiotti, R Bajpai, M A Barton, K Cannon, E Capocasa, M Chan, C Chen, K Chen, Y Chen, H Chu, Y K Chu, S Eguchi, Y Enomoto, R Flaminio, Y Fujii, M Fukunaga, M Fukushima, G Ge, A Hagiwara, S Haino, K Hasegawa, H Hayakawa, K Hayama, Y Himemoto, Y Hiranuma, N Hirata, E Hirose, Z Hong, B H Hsieh, C Z Huang, P Huang, Y Huang, B Ikenoue, S Imam, K Inayoshi, Y Inoue, K Ioka, Y Itoh, K Izumi, K Jung, P Jung, T Kajita, M Kamiizumi, N Kanda, G Kang, K Kawaguchi, N Kawai, T Kawasaki, C Kim, J C Kim, W S Kim, Y M Kim, N Kimura, N Kita, H Kitazawa, Y Kojima, K Kokeyama, K Komori, A K H Kong, K Kotake, C Kozakai, R Kozu, R Kumar, J Kume, C Kuo, H S Kuo, S Kuroyanagi, K Kusayanagi, K Kwak, H K Lee, H W Lee, R Lee, M Leonardi, L C C Lin, C Y Lin, F L Lin, G C Liu, L W Luo, M Marchio, Y Michimura, N Mio, O Miyakawa, A Miyamoto, Y Miyazaki, K Miyo, S Miyoki, S Morisaki, Y Moriwaki, K Nagano, S Nagano, K Nakamura, H Nakano, M Nakano, R Nakashima, T Narikawa, R Negishi, W T Ni, A Nishizawa, Y Obuchi, W Ogaki, J J Oh, S H Oh, M Ohashi, N Ohishi, M Ohkawa, K Okutomi, K Oohara, C P Ooi, S Oshino, K Pan, H Pang, J Park, F E Peña Arellano, I Pinto, N Sago, S Saito, Y Saito, K Sakai, Y Sakai, Y Sakuno, S Sato, T Sato, T Sawada, T Sekiguchi, Y Sekiguchi, S Shibagaki, R Shimizu, T Shimoda, K Shimode, H Shinkai, T Shishido, A Shoda, K Somiya, E J Son, H Sotani, R Sugimoto, T Suzuki, T Suzuki, H Tagoshi, H Takahashi, R Takahashi, A Takamori, S Takano, H Takeda, M Takeda, H Tanaka, K Tanaka, K Tanaka, T Tanaka, T Tanaka, S Tanioka, E N Tapia San Martin, S Telada, T Tomaru, Y Tomigami, T Tomura, F Travasso, L Trozzo, T Tsang, K Tsubono, S Tsuchida, T Tsuzuki, D Tuyenbayev, N Uchikata, T Uchiyama, A Ueda, T Uehara, K Ueno, G Ueshima, F Uraguchi, T Ushiba, M H P M van Putten, H Vocca, J Wang, C Wu, H Wu, S Wu, W-R Xu, T Yamada, K Yamamoto, K Yamamoto, T Yamamoto, K Yokogawa, J Yokoyama, T Yokozawa, T Yoshioka, H Yuzurihara, S Zeidler, Y Zhao, and Z H Zhu. Overview of kagra: Detector design and construction history. *Progress of Theoretical and Experimental Physics*, 2021(5):05A101, 08 2020.

[53] Matthew Evans, Rana X Adhikari, Chaitanya Afle, Stefan W. Ballmer, Sylvia Biscoveanu, Ssohrab Borhanian, Duncan A. Brown, Yanbei Chen, Robert Eisenstein, Alexandra Gruson, Anuradha Gupta, Evan D. Hall, Rachael Huxford, Brittany Kamai, Rahul Kashyap, Jeff S. Kissel, Kevin Kuns, Philippe Landry, Amber Lenon, Geoffrey Lovelace, Lee McCuller,

Ken K. Y. Ng, Alexander H. Nitz, Jocelyn Read, B. S. Sathyaprakash, David H. Shoemaker, Bram J. J. Slagmolen, Joshua R. Smith, Varun Srivastava, Ling Sun, Salvatore Vitale, and Rainer Weiss. A horizon study for cosmic explorer: Science, observatories, and community, 2021. *arXiv e-Print Archive*, arXiv:2109.09882 [astro-ph.IM].

[54] M Punturo, M Abernathy, F Acernese, B Allen, N Andersson, K Arun, F Barone, B Barr, M Barsuglia, M Beker, N Beveridge, S Birindelli, S Bose, L Bosi, S Braccini, C Bradaschia, T Bulik, E Calloni, G Cella, E Chassande Mottin, S Chelkowski, A Chincarini, J Clark, E Coccia, C Colacino, J Colas, A Cumming, L Cunningham, E Cuoco, S Danilishin, K Danzmann, G De Luca, R De Salvo, T Dent, R Derosa, L Di Fiore, A Di Virgilio, M Doets, V Fafone, P Falferi, R Flaminio, J Franc, F Frasconi, A Freise, P Fulda, J Gair, G Gemme, A Gennai, A Giazotto, K Glampedakis, M Granata, H Grote, G Guidi, G Hammond, M Hannam, J Harms, D Heinert, M Hendry, I Heng, E Hennes, S Hild, J Hough, S Husa, S Huttner, G Jones, F Khalili, K Kokeyama, K Kokkotas, B Krishnan, M Lorenzini, H Lück, E Majorana, I Mandel, V Mandic, I Martin, C Michel, Y Minenkov, N Morgado, S Mosca, B Mours, H Müller-Ebhardt, P Murray, R Nawrodt, J Nelson, R Oshaughnessy, C D Ott, C Palomba, A Paoli, G Parguez, A Pasqualetti, R Passaquieti, D Passuello, L Pinard, R Poggiani, P Popolizio, M Prato, P Puppo, D Rabeling, P Rapagnani, J Read, T Regimbau, H Rehbein, S Reid, L Rezzolla, F Ricci, F Richard, A Rocchi, S Rowan, A Rüdiger, B Sassolas, B Sathyaprakash, R Schnabel, C Schwarz, P Seidel, A Sintes, K Somiya, F Speirits, K Strain, S Strigin, P Sutton, S Tarabrin, J van den Brand, C van Leewen, M van Veggel, C van den Broeck, A Vecchio, J Veitch, F Vetrano, A Vicere, S Vyatchanin, B Willke, G Woan, P Wolfango, and K Yamamoto. The third generation of gravitational wave observatories and their science reach. *Classical and Quantum Gravity*, 27(8):084007, apr 2010.

[55] M. Maggiore et al. Science case for the einstein telescope. *Journal of Cosmology and Astroparticle Physics*, 2020(03):050, mar 2020.

[56] Stanislav Babak, Martin Hewitson, and Antoine Petiteau. Lisa sensitivity and snr calculations, 2021. *arXiv e-Print Archive*, arXiv:2108.01167 [astro-ph.IM].

[57] J. Baker et al. The laser interferometer space antenna: Unveiling the millihertz gravitational wave sky, 2019. *arXiv e-Print Archive*, arXiv:1907.06482.

[58] F. Jenet, L. S. Finn, J. Lazio, A. Lommen, M. McLaughlin, I. Stairs, D. Stinebring, J. Verbiest, A. Archibald, Z. Arzoumanian, D. Backer, J. Cordes, P. Demorest, R. Ferdman, P. Freire, M. Gonzalez, V. Kaspi, V. Kondratiev, D. Lorimer, R. Lynch, D. Nice, S. Ransom, R. Shannon, and X. Siemens. The north american nanohertz observatory for gravitational waves, 2009. *arXiv e-Print Archive*, arXiv:0909.1058 [astro-ph.IM].

[59] Péter Mészáros, Derek B. Fox, Chad Hanna, and Kohta Murase. Multimessenger astrophysics. *Nature Reviews Physics*, 1(10):585–599, Oct 2019.

[60] Karl Schwarzschild. Über das gravitationsfeld eines massenpunktes nach der einsteinschen theorie.

[61] P. Renteln. *Manifolds, Tensors and Forms.* Cambridge University Press, 2014.

[62] Daniel Grumiller and Mohammad Mehdi Sheikh-Jabbari. *Black Hole Physics: From Collapse to Evaporation.* Grad.Texts Math. Springer, 11 2022.

[63] Roy P. Kerr. Gravitational field of a spinning mass as an example of algebraically special metrics. *Phys. Rev. Lett.*, 11:237–238, Sep 1963.

[64] Stephen Hawking. Gravitationally Collapsed Objects of Very Low Mass. *Monthly Notices of the Royal Astronomical Society*, 152(1):75–78, 04 1971.

[65] Bernard Carr and Steven Giddings. Quantum black holes. *Scientific American*, 292:48–55, 06 2005.

[66] I. Bombaci. The maximum mass of a neutron star. *Astronomy and Astrophysics*, 305:871, January 1996.

[67] Luciano Rezzolla, Elias R. Most, and Lukas R. Weih. Using gravitational-wave observations and quasi-universal relations to constrain the maximum mass of neutron stars. *The Astrophysical Journal Letters*, 852(2):L25, January 2018.

[68] G. Brown and H. A. Bethe. How a supernova explodes. *Scientific American*, 252:60–68, May 1985.

[69] B. J. Carr and S. W. Hawking. Black Holes in the Early Universe. *Monthly Notices of the Royal Astronomical Society*, 168(2):399–415, 08 1974.

[70] B. Paczyński. Evolutionary Processes in Close Binary Systems. *Annual Review of Astronomy and Astrophysics*, 9:183, January 1971.

[71] Theocharis A. Apostolatos, Curt Cutler, Gerald J. Sussman, and Kip S. Thorne. Spin-induced orbital precession and its modulation of the gravitational waveforms from merging binaries. *Phys. Rev. D*, 49:6274–6297, Jun 1994.

[72] Curt Cutler and Éanna E. Flanagan. Gravitational waves from merging compact binaries: How accurately can one extract the binary's parameters from the inspiral waveform? *Physical Review D*, 49(6):2658–2697, March 1994.

[73] Patricia Schmidt, Frank Ohme, and Mark Hannam. Towards models of gravitational waveforms from generic binaries: Ii. modelling precession effects with a single effective precession parameter. *Phys. Rev. D*, 91:024043, Jan 2015.

[74] K. Atkinson and W. Han. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*. Lecture Notes in Mathematics. Springer, 2012.

[75] Y. Kosmann-Schwarzbach and S.F. Singer. *Groups and Symmetries: From Finite Groups to Lie Groups*. Universitext. Springer New York, 2009.

[76] Ian Harry, Juan Calderón Bustillo, and Alex Nitz. Searching for the full symphony of black hole binary mergers. *Phys. Rev. D*, 97:023004, Jan 2018.

[77] J. N. Goldberg, A. J. Macfarlane, E. T. Newman, F. Rohrlich, and E. C. G. Sudarshan. Spin-s Spherical Harmonics and Thop. *Journal of Mathematical Physics*, 8(11):2155–2161, November 1967.

[78] T.W. Baumgarte and S.L. Shapiro. *Numerical Relativity: Solving Einstein's Equations on the Computer*. Cambridge University Press, 2010.

[79] Frans Pretorius. Evolution of binary black-hole spacetimes. *Phys. Rev. Lett.*, 95:121101, Sep 2005.

[80] Karan Jani, James Healy, James A Clark, Lionel London, Pablo Laguna, and Deirdre Shoemaker. Georgia tech catalog of gravitational waveforms. *Classical and Quantum Gravity*, 33(20):204001, September 2016.

[81] James Healy, Carlos O Lousto, Yosef Zlochower, and Manuela Campanelli. The rit binary black hole simulations catalog. *Classical and Quantum Gravity*, 34(22):224001, oct 2017.

[82] James Healy and Carlos O. Lousto. Fourth rit binary black hole simulations catalog: Extension to eccentric orbits. *Physical Review D*, 105(12), June 2022.

[83] Michael Pürrer. Frequency-domain reduced order models for gravitational waves from aligned-spin compact binaries. *Classical and Quantum Gravity*, 31(19):195010, September 2014.

[84] Scott E. Field, Chad R. Galley, Jan S. Hesthaven, Jason Kaye, and Manuel Tiglio. Fast prediction and evaluation of gravitational waveforms using surrogate models. *Physical Review X*, 4(3), July 2014.

[85] Vijay Varma, Scott E. Field, Mark A. Scheel, Jonathan Blackman, Lawrence E. Kidder, and Harald P. Pfeiffer. Surrogate model of hybridized numerical relativity binary black hole waveforms. *Physical Review D*, 99(6), March 2019.

[86] Alessandra Buonanno, Bala R. Iyer, Evan Ochsner, Yi Pan, and B. S. Sathyaprakash. Comparison of post-newtonian templates for compact binary inspiral signals in gravitational-wave detectors. *Phys. Rev. D*, 80:084043, Oct 2009.

[87] Francesco Messina, Reetika Dudi, Alessandro Nagar, and Sebastiano Bernuzzi. Quasi-5.5pn taylorf2 approximant for compact binaries: Point-mass phasing and impact on the tidal polarizability inference. *Phys. Rev. D*, 99:124051, Jun 2019.

[88] A. Buonanno and T. Damour. Effective one-body approach to general relativistic two-body dynamics. *Phys. Rev. D*, 59:084006, Mar 1999.

[89] Rossella Gamba, Sarp Akçay, Sebastiano Bernuzzi, and Jake Williams. Effective-one-body waveforms for precessing coalescing compact binaries with post-newtonian twist. *Phys. Rev. D*, 106:024020, Jul 2022.

[90] Geraint Pratten, Cecilio García-Quirós, Marta Colleoni, Antoni Ramos-Buades, Héctor Estellés, Maite Mateu-Lucena, Rafel Jaume, Maria Haney, David Keitel, Jonathan E. Thompson, and Sascha Husa. Computationally efficient models for the dominant and subdominant harmonic modes of precessing binary black holes. *Phys. Rev. D*, 103:104056, May 2021.

[91] Cecilio García-Quirós, Marta Colleoni, Sascha Husa, Héctor Estellés, Geraint Pratten, Antoni Ramos-Buades, Maite Mateu-Lucena, and Rafel Jaume. Multimode frequency-domain model for the gravitational wave signal from nonprecessing black-hole binaries. *Phys. Rev. D*, 102:064002, Sep 2020.

[92] Geraint Pratten, Sascha Husa, Cecilio García-Quirós, Marta Colleoni, Antoni Ramos-Buades, Héctor Estellés, and Rafel Jaume. Setting the cornerstone for a family of models for gravitational waves from compact binaries: The dominant harmonic for nonprecessing quasicircular black holes. *Phys. Rev. D*, 102:064001, Sep 2020.

[93] Héctor Estellés, Antoni Ramos-Buades, Sascha Husa, Cecilio García-Quirós, Marta Colleoni, Leïla Haegel, and Rafel Jaume. Phenomenological time domain model for dominant quadrupole gravitational wave signal of coalescing binary black holes. *Phys. Rev. D*, 103:124060, Jun 2021.

[94] Héctor Estellés, Sascha Husa, Marta Colleoni, David Keitel, Maite Mateu-Lucena, Cecilio García-Quirós, Antoni Ramos-Buades, and Angela Borchers. Time domain phenomenological model of gravitational wave subdominant harmonics for quasi-circular non-precessing binary black hole coalescences, 2020. *arXiv e-Print Archive*, arXiv:2012.11923 [gr-qc].

[95] Sebastian Khan, Sascha Husa, Mark Hannam, Frank Ohme, Michael Pürrer, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. ii. a phenomenological

model for the advanced detector era. *Physical Review D*, 93(4), February 2016.

[96] L.D. Landau, E.M. Lifshitz, and J.B. Sykes. *Mechanics: Volume 1*. Course of theoretical physics. Elsevier Science, 1976.

[97] B.C. Hall. *Quantum Theory for Mathematicians*. Graduate Texts in Mathematics. Springer New York, 2013.

[98] G.B. Folland. *Quantum Field Theory: A Tourist Guide for Mathematicians*. Mathematical surveys and monographs. American Mathematical Society, 2008.

[99] T. Lancaster and S. Blundell. *Quantum Field Theory for the Gifted Amateur*. OUP Oxford, 2014.

[100] M.E. Peskin. *An Introduction To Quantum Field Theory*. CRC Press, 2018.

[101] A. Vilenkin, A. Vilenkin, and E.P.S. Shellard. *Cosmic Strings and Other Topological Defects*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 1994.

[102] J.R. Munkres. *Topology*. Featured Titles for Topology. Prentice Hall, Incorporated, 2000.

[103] G.E. Bredon. *Topology and Geometry*. Number v. 14 in Graduate Texts in Mathematics. Springer, 1993.

[104] A. Hatcher. *Algebraic Topology*. Algebraic Topology. Cambridge University Press, 2002.

[105] M. Spivak. *Physics for Mathematicians: Mechanics I*. Physics for Mathematicians: Mechanics I. Publish or Perish, 2010.

[106] K. Vogtmann, A. Weinstein, and V.I. Arnol'd. *Mathematical Methods of Classical Mechanics*. Graduate Texts in Mathematics. Springer New York, 1997.

[107] D.J. Griffiths. *Introduction to Electrodynamics*. Pearson Education, 2014.

[108] D.J. Griffiths. *Introduction to Quantum Mechanics*. Cambridge University Press, 2017.

[109] M.A. Nielsen and I.L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[110] Stephen W. Hawking. Black hole explosions? *Nature*, 248:30–31, 1974.

[111] S. W. Hawking. Breakdown of predictability in gravitational collapse. *Phys. Rev. D*, 14:2460–2473, Nov 1976.

[112] Michael F. Atiyah. Topological quantum field theory. *Publications Mathématiques de l'IHÉS*, 68:175–186, 1988.

[113] Eric C. Rowell and Zhenghan Wang. Mathematics of topological quantum computing, 2017. *arXiv e-Print Archive*, arXiv:1705.06206 [math.QA].

[114] R.F. Streater and A.S. Wightman. *PCT, Spin and Statistics, and All that*. Princeton landmarks in mathematics and physics. Princeton University Press, 2000.

[115] M. Vázquez and A. Hanslmeier. *Ultraviolet Radiation in the Solar System*. Astrophysics and Space Science Library. Springer Netherlands, 2005.

[116] I.M. Gel'fand and G.E. Shilov. *Generalized Functions: Properties and operations, by I. M. Gel'fand and G. E. Shilov, translated by E. Saletan*. Generalized Functions. Academic Press, 1964.

[117] B. Zwiebach. *A First Course in String Theory*. A First Course in String Theory. Cambridge University Press, 2004.

[118] P. Deligne and N.J.) Institute for Advanced Study (Princeton. *Quantum Fields and Strings: A Course for Mathematicians*. Number v. 1 in Quantum Fields and Strings: A Course for Mathematicians. American Mathematical Society, 1999.

[119] Juan Maldacena. *International Journal of Theoretical Physics*, 38(4):1113–1133, 1999.

[120] Mairi Sakellariadou. Cosmic strings and cosmic superstrings. *Nuclear Physics B - Proceedings Supplements*, 192-193:68–90, 2009. Theory and Particle Physics: The LHC Perspective and Beyond.

[121] E. Yankowsky, A.S. Schwarz, and S. Levy. *Quantum Field Theory and Topology*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1993.

[122] Edmund J. Copeland and T. W. B. Kibble. Cosmic strings and superstrings. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2115):623–657, January 2010.

[123] Jose J. Blanco-Pillado, Ken D. Olum, and Benjamin Shlaer. Number of cosmic string loops. *Phys. Rev. D*, 89:023512, Jan 2014.

[124] Larissa Lorenz, Christophe Ringeval, and Mairi Sakellariadou. Cosmic string loop distribution on all length scales and at any redshift. *Journal of Cosmology and Astroparticle Physics*, 2010(10):003, oct 2010.

[125] Andreas Albrecht and Neil Turok. Evolution of cosmic string networks. *Phys. Rev. D*, 40:973–1001, Aug 1989.

[126] David P. Bennett and François R. Bouchet. High-resolution simulations of cosmic-string evolution. i. network evolution. *Phys. Rev. D*, 41:2408–2433, Apr 1990.

[127] B. Allen and E. P. S. Shellard. Cosmic-string evolution: A numerical simulation. *Phys. Rev. Lett.*, 64:119–122, Jan 1990.

[128] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, M. Agathos, K. Agatsuma, N. Aggarwal, and O. D. Aguiar. Constraints on cosmic strings using data from the third advanced ligo–virgo observing run. *Phys. Rev. Lett.*, 126:241102, Jun 2021.

[129] E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, 2010.

[130] Judith C. Brown. Calculation of a constant Q spectral transform. *Acoustical Society of America Journal*, 89(1):425–434, January 1991.

[131] M. Hall. Resolution and uncertainty in spectral decomposition. *First Break*, 24(12), 2006.

[132] C. Chatfield. *The Analysis of Time Series: An Introduction, Sixth Edition*. Chapman & Hall/CRC Texts in Statistical Science. CRC Press, 2016.

[133] D Davis, J S Areeda, B K Berger, R Bruntz, A Effler, R C Essick, R P Fisher, P Godwin, E Goetz, A F Helmling-Cornell, B Hughey, E Katsavounidis, A P Lundgren, D M Macleod, Z Márka, T J Massinger, A Matas, J McIver, G Mo, K Mogushi, and P Nguyen. Ligo detector characterization in the second and third observing runs. *Classical and Quantum Gravity*, 38(13):135014, June 2021.

[134] C.W. Helstrom. Statistical theory of signal detection. 1960.

[135] B. S. Sathyaprakash and S. V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 44:3819–3834, Dec 1991.

[136] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction To Algorithms*. Mit Electrical Engineering and Computer Science. MIT Press, 2001.

[137] Bruce Allen. A chi-squared time-frequency discriminator for gravitational wave detection. *Phys. Rev. D*, 71:062001, Mar 2005.

[138] J. Veitch and A. Vecchio. Bayesian coherent analysis of in-spiral gravitational wave signals with a detector network. *Physical Review D*, 81(6), March 2010.

[139] Nelson Christensen and Renate Meyer. Parameter estimation with gravitational waves. *Rev. Mod. Phys.*, 94:025001, Apr 2022.

[140] Rich Abbott et al. Open data from the first and second observing runs of Advanced LIGO and Advanced Virgo. *SoftwareX*, 13:100658, 2021.

[141] R. Abbott et al. Open Data from the Third Observing Run of LIGO, Virgo, KAGRA, and GEO. *Astrophys. J. Suppl.*, 267(2):29, 2023.

[142] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, G. Allen, A. Allocca, M. A. Aloy, P. A. Altin, A. Amato, A. Ananyeva, S. B. Anderson, W. G. Anderson, S. V. Angelova, S. Antier, S. Appert, and K. Arai. Gwtc-1: A gravitational-wave transient catalog of compact binary mergers observed by ligo and virgo during the first and second observing runs. *Phys. Rev. X*, 9:031040, Sep 2019.

[143] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, S. Akcay, G. Allen, A. Allocca, P. A. Altin, A. Amato, S. Anand, A. Ananyeva, S. B. Anderson, W. G. Anderson, S. V. Angelova, S. Ansoldi, J. M. Antelis, S. Antier, S. Appert, K. Arai, M. C. Araya, J. S. Areeda, M. Arène, N. Arnaud, S. M. Aronson, K. G. Arun, Y. Asali, S. Ascenzi, and G. Ashton. Gwtc-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *Phys. Rev. X*, 11:021053, Jun 2021.

[144] The LIGO Scientific Collaboration, the Virgo Collaboration, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, S. Albanesi, A. Allocca, P. A. Altin, A. Amato, C. Anand, S. Anand, A. Ananyeva, S. B. Anderson, and W. G. Anderson. Gwtc-2.1: Deep extended catalog of compact binary coalescences observed by ligo and virgo during the first half of the third observing run, 2022. *arXiv e-Print Archive*, arXiv:2108.01045 [gr-qc].

[145] LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite. free software (GPL), 2018.

[146] D. M. Macleod, J. S. Areeda, S. B. Coughlin, T. J. Massinger, and A. L. Urban. GWpy: A Python package for gravitational-wave astrophysics. *SoftwareX*, 13:100657, 2021.

[147] Alex Nitz, Ian Harry, Duncan Brown, Christopher M. Biwer, Josh Willis, Tito Dal Canton, Collin Capano, Thomas Dent, Larne Pekowsky, Gareth S Cabourn Davies, Soumi De, Miriam Cabero, Shichao Wu, Andrew R. Williamson, Bernd Machenschalk, Duncan Macleod, Francesco Pannarale, Prayush Kumar, Steven Reyes, dfinstad, Sumit Kumar, Márton Tápai, Leo Singer, Praveen Kumar, veronica villa, maxtrevor, Bhooshan

Uday Varsha Gadre, Sebastian Khan, Stephen Fairhurst, and Arthur Tolley. gwastro/pycbc: v2.3.3 release of pycbc, January 2024.

[148] Samantha A Usman, Alexander H Nitz, Ian W Harry, Christopher M Biwer, Duncan A Brown, Miriam Cabero, Collin D Capano, Tito Dal Canton, Thomas Dent, Stephen Fairhurst, Marcel S Kehl, Drew Keppel, Badri Krishnan, Amber Lenon, Andrew Lundgren, Alex B Nielsen, Larne P Pekowsky, Harald P Pfeiffer, Peter R Saulson, Matthew West, and Joshua L Willis. The pycbc search for gravitational waves from compact binary coalescence. *Classical and Quantum Gravity*, 33(21):215004, October 2016.

[149] Kipp Cannon, Sarah Caudill, Chiwai Chan, Bryce Cousins, Jolien D.E. Creighton, Becca Ewing, Heather Fong, Patrick Godwin, Chad Hanna, Shaun Hooper, Rachael Huxford, Ryan Magee, Duncan Meacher, Cody Messick, Soichiro Morisaki, Debnandini Mukherjee, Hiroaki Ohta, Alexander Pace, Stephen Privitera, Iris de Ruiter, Surabhi Sachdev, Leo Singer, Divya Singh, Ron Tapia, Leo Tsukada, Daichi Tsuna, Takuya Tsutsui, Koh Ueno, Aaron Viets, Leslie Wade, and Madeline Wade. Gstlal: A software framework for gravitational wave discovery. *SoftwareX*, 14:100680, 2021.

[150] Wim Taymans, Steve Baker, Andy Wingo, Ronald S. Bultje, and Stefan Kost. *GStreamer 1.8.3 Application Development Manual.* Samurai Media Limited, London, GBR, 2016.

[151] Neil J Cornish and Tyson B Littenberg. Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches. *Classical and Quantum Gravity*, 32(13):135012, June 2015.

[152] Neil J. Cornish, Tyson B. Littenberg, Bence Bécsy, Katerina Chatziioannou, James A. Clark, Sudarshan Ghonge, and Margaret Millhouse. Bayeswave analysis pipeline in the era of gravitational wave observations. *Physical Review D*, 103(4), February 2021.

[153] S. Mallat. *A Wavelet Tour of Signal Processing: The Sparse Way.* Elsevier Science, 2008.

[154] Marco Drago, Sergey Klimenko, Claudia Lazzaro, Edoardo Milotti, Guenakh Mitselmakher, Valentin Necula, Brendan O'Brian, Giovanni Andrea Prodi, Francesco Salemi, Marek Szczepanczyk, Shubhanshu Tiwari, Vaibhav Tiwari, Gayathri V, Gabriele Vedovato, and Igor Yakushin. coherent waveburst, a pipeline for unmodeled gravitational-wave data analysis. *SoftwareX*, 14:100678, June 2021.

[155] Gregory Ashton, Moritz Hübner, Paul D. Lasky, Colm Talbot, Kendall Ackley, Sylvia Biscoveanu, Qi Chu, Atul Divakarla, Paul J. Easter, Boris Goncharov, Francisco Hernandez Vivanco, Jan Harms, Marcus E. Lower, Grant D. Meadors, Denyz Melchor, Ethan Payne, Matthew D. Pitkin, Jade

Powell, Nikhil Sarin, Rory J. E. Smith, and Eric Thrane. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(2):27, April 2019.

[156] G. James, D. Witten, T. Hastie, and R. Tibshirani. *An Introduction to Statistical Learning: with Applications in R*. Springer Texts in Statistics. Springer New York, 2014.

[157] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*. Springer Series in Statistics. Springer New York, 2009.

[158] S.J. Russell, P. Norvig, and E. Davis. *Artificial Intelligence: A Modern Approach*. Prentice Hall series in artificial intelligence. Prentice Hall, 2010.

[159] C.M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer, 2006.

[160] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. Adaptive Computation and Machine Learning series. MIT Press, 2016.

[161] Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biology*, 52:99–115, 1990.

[162] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386–408, 1958.

[163] A. Bondy and U.S.R. Murty. *Graph Theory*. Graduate Texts in Mathematics. Springer London, 2011.

[164] Avrim L. Blum and Ronald L. Rivest. Training a 3-node neural network is np-complete. *Neural Networks*, 5(1):117–127, 1992.

[165] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. *arXiv e-Print Archive*, arXiv:1412.6980 [cs.LG].

[166] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(61):2121–2159, 2011.

[167] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.

[168] Leonardo Ferreira Guilhoto. An overview of artificial neural networks for mathematicians. 2018.

[169] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[170] George V. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2:303–314, 1989.

[171] W. Rudin. *Principles of Mathematical Analysis*. International series in pure and applied mathematics. McGraw-Hill, 1976.

[172] Halbert White. Connectionist nonparametric regression: multilayer feedforward networks can learn arbitrary mappings. *Neural Netw.*, 3(5):535–549, October 1990.

[173] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251–257, 1991.

[174] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. *arXiv e-Print Archive*, arXiv:1912.01703 [cs.LG].

[175] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

[176] NVIDIA, Péter Vingelmann, and Frank H.P. Fitzek. Cuda, release: 10.2.89, 2020.

[177] Jayshree Ghorpade. Gpgpu processing in cuda architecture. *Advanced Computing: An International Journal*, 3(1):105–120, January 2012.

[178] Vaishak Belle and Ioannis Papantonis. Principles and practice of explainable machine learning, 2020. *arXiv e-Print Archive*, arXiv:2009.11698 [cs.LG].

[179] Pradnya A. Vikhar. Evolutionary algorithms: A critical review and its future prospects. In *2016 International Conference on Global Trends in Signal Processing, Information Computing and Communication (ICGTSPICC)*, pages 261–265, 2016.

[180] Lee Altenberg. The schema theorem and price's theorem. volume 3 of *Foundations of Genetic Algorithms*, pages 23–49. Elsevier, 1995.

[181] John A. Lee and Michel Verleysen. *Nonlinear Dimensionality Reduction.* Springer Publishing Company, Incorporated, 1st edition, 2007.

[182] Marina Meilă and Hanyu Zhang. Manifold Learning: What, How, and Why. *Annual Review of Statistics and Its Application*, 11(1):annurev, March 2024.

[183] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008.

[184] Geoffrey E Hinton and Sam Roweis. Stochastic neighbor embedding. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems*, volume 15. MIT Press, 2002.

[185] T.M. Cover and J.A. Thomas. *Elements of Information Theory.* Wiley, 2012.

[186] J.A. Rice. *Mathematical Statistics and Data Analysis.* Advanced series. Cengage Learning, 2007.

[187] Joseph Lee Rodgers and W. Alan Nicewander. Thirteen ways to look at the correlation coefficient. *The American Statistician*, 42(1):59–66, 1988.

[188] T. Damour and A. Vilenkin. Gravitational Radiation from Cosmic (Super)strings: Bursts, Stochastic Background, and Observational Windows. *Physical Review D*, 71(6):063510, March 2005.

[189] T. Damour and A. Vilenkin. Gravitational Wave Bursts from Cosmic Strings. *Physical Review Letters*, 85(18):3761–3764, October 2000.

[190] T. Damour and A. Vilenkin. Gravitational Wave Bursts from Cusps and Kinks on Cosmic Strings. *Physical Review D*, 64(6):064008, September 2001.

[191] X. Siemens, J. Creighton, I. Maor, S.R. Majumder, K. Cannon, and J. Read. Gravitational wave bursts from cosmic (super)strings: Quantitative analysis and constraints. *Phys. Rev. D*, 73:105001, May 2006.

[192] J. Aasi et al. Constraints on cosmic strings from the LIGO-virgo gravitational-wave detectors. *Physical Review Letters*, 112(13), apr 2014.

[193] B.P. Abbott et al. First LIGO search for gravitational wave bursts from cosmic (super)strings. *Physical Review D*, 80(6), sep 2009.

[194] B.P. Abbott et al. Constraints on cosmic strings using data from the first advanced LIGO observing run. *Physical Review D*, 97(10), may 2018.

[195] M. Cabero et al. Blip glitches in advanced ligo data. *Classical and Quantum Gravity*, 36(15):155010, jul 2019.

[196] J. Yan, M. Avagyan, R.E. Colgan, D. Veske, I. Bartos, J. Wright, Z. Marka, and S. Marka. Generalized approach to matched filtering using neural networks. *Phys. Rev. D*, 105:043006, Feb 2022.

[197] ET Steering Committee Editorial Team. Design report update 2020 for the einstein telescope. Technical report, 2020.

[198] Unofficial sensitivity curves (asd) for aligo, kagra, virgo, voyager, cosmic explorer, and einstein telescope. `https://dcc.ligo.org/LIGO-T1500293/public`, 2020.

[199] Noise curves used for simulations in the update of the observing scenarios paper. `https://dcc.ligo.org/LIGO-T2000012/public`, 2022.

[200] M. Lopez and S. Schmidt. Documentation of the gengli Package. `https://melissa.lopez.docs.ligo.org/gengli/index.html`, 2022.

[201] M. Lopez, V. Boudart, K. Buijsman, A. Reza, and S. Caudill. Simulating Transient Noise Bursts in LIGO with Generative Adversarial Networks. *Phys. Rev. D*, 106(2):023027, 2022.

[202] M. Lopez et al. Simulating Transient Noise Bursts in LIGO with gengli. 5 2022. *arXiv e-Print Archive*, arXiv:2205.09204 [astro-ph.IM].

[203] A. van den Oord et al. WaveNet: A Generative Model for Raw Audio. *arXiv e-prints*, page arXiv:1609.03499, September 2016.

[204] A. van den Oord et al. Conditional Image Generation with PixelCNN Decoders. *arXiv e-prints*, page arXiv:1606.05328, June 2016.

[205] C. Szegedy et al. Going Deeper with Convolutions. *arXiv e-prints*, page arXiv:1409.4842, September 2014.

[206] K. He et al. Deep Residual Learning for Image Recognition. *arXiv e-prints*, page arXiv:1512.03385, December 2015.

[207] W. Wei et al. Deep Learning Ensemble for Real-time Gravitational Wave Detection of Spinning Binary Black Hole Mergers. *Phys. Lett. B*, 812:136029, 2021.

[208] I. Loshchilov and F. Hutter. Decoupled Weight Decay Regularization. *arXiv e-prints*, page arXiv:1711.05101, November 2017.

[209] T. Kadir and M. Brady. Saliency, scale and image description. *International Journal of Computer Vision*, 45(2):83–105, 2001.

[210] N. Kokhlikyan et al. Captum: A unified and generic model interpretability library for pytorch. *CoRR*, abs/2009.07896, 2020.

[211] I. Jolliffe and J. Cadima. Principal component analysis: A review and recent developments. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374:20150202, 04 2016.

[212] S. Gillessen, F. Eisenhauer, S. Trippe, T. Alexander, R. Genzel, F. Martins, and T. Ott. Monitoring stellar orbits around the massive black hole in the galactic center. *The Astrophysical Journal*, 692(2):1075, feb 2009.

[213] A. M. Ghez, B. L. Klein, M. Morris, and E. E. Becklin. High proper-motion stars in the vicinity of sagittarius a*: Evidence for a supermassive black hole at the center of our galaxy. *The Astrophysical Journal*, 509(2):678, dec 1998.

[214] S. Bowyer, E. T. Byram, T. A. Chubb, and H. Friedman. Cosmic x-ray sources. *Science*, 147(3656):394–398, 1965.

[215] W. M. Goss and R. X. McGee. The Discovery of the Radio Source Sagittarius A (Sgr A). In Roland Gredel, editor, *The Galactic Center*, volume 102 of *Astronomical Society of the Pacific Conference Series*, page 369, January 1996.

[216] R. Genzel, D. Hollenbach, and C. H. Townes. The nucleus of our Galaxy. *Reports on Progress in Physics*, 57(5):417–479, May 1994.

[217] H.L. Shipman. The Implausible History of Triple Star Models for Cygnus X-1: Evidence for a Black Hole. *Astrophysical Letters*, 16:9, February 1975.

[218] LIGO Scientific Collaboration and Virgo Collaboration. Gw190814: Gravitational waves from the coalescence of a 23 solar mass black hole with a 2.6 solar mass compact object. *The Astrophysical Journal Letters*, 896(2):L44, jun 2020.

[219] Jenny E. Greene, Jay Strader, and Luis C. Ho. Intermediate-mass black holes. *Annual Review of Astronomy and Astrophysics*, 58(1):257–312, August 2020.

[220] R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, N. Adhikari, R. X. Adhikari, V. B. Adya, C. Affeldt, D. Agarwal, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, and T. Akutsu. Search for intermediate-mass black hole binaries in the third observing run of advanced ligo and advanced virgo. *Astronomy & Astrophysics*, 659:A84, March 2022.

[221] Zoltán Haiman and Abraham Loeb. What Is the Highest Plausible Redshift of Luminous Quasars? *The Astrophysics Journal*, 552(2):459–463, May 2001.

[222] Marta Volonteri, Francesco Haardt, and Piero Madau. The Assembly and Merging History of Supermassive Black Holes in Hierarchical Models of Galaxy Formation. *The Astrophysics Journal*, 582(2):559–573, January 2003.

[223] Davide Gerosa and Maya Fishbach. Hierarchical mergers of stellar-mass black holes and their gravitational-wave signatures. *Nature Astronomy*, 5(8):749–760, July 2021.

[224] Gulab C. Dewangan, Lev Titarchuk, and Richard E. Griffiths. Black hole mass of the ultraluminous x-ray source m82 x-1. *The Astrophysical Journal*, 637(1):L21, jan 2006.

[225] N. Madhusudhan, S. Justham, L. Nelson, B. Paxton, E. Pfahl, Ph. Podsiadlowski, and S. Rappaport. Models of ultraluminous x-ray sources with intermediate-mass black holes. *The Astrophysical Journal*, 640(2):918, apr 2006.

[226] Sayantan Ghosh, Koustav Chandra, and Archana Pai. Unmasking noise transients masquerading as intermediate-mass black hole binaries, 2023.

[227] Koustav Chandra, Juan Calderón Bustillo, Archana Pai, and I.W. Harry. First gravitational-wave search for intermediate-mass black hole mergers with higher-order harmonics. *Physical Review D*, 106(12), December 2022.

[228] He Wang, Shichao Wu, Zhoujian Cao, Xiaolin Liu, and Jian-Yang Zhu. Gravitational-wave signal recognition of ligo data by deep learning. *Phys. Rev. D*, 101:104003, May 2020.

[229] Daniel George and E. A. Huerta. Deep learning for real-time gravitational wave detection and parameter estimation with ligo data, 2017. *arXiv e-Print Archive*, arXiv:1711.07966 [gr-qc].

[230] Daniel George and E.A. Huerta. Deep learning for real-time gravitational wave detection and parameter estimation: Results with advanced ligo data. *Physics Letters B*, 778:64–70, 2018.

[231] Timothy D. Gebhard, Niki Kilbertus, Ian Harry, and Bernhard Schölkopf. Convolutional neural networks: A magic bullet for gravitational-wave detection? *Phys. Rev. D*, 100:063015, Sep 2019.

[232] XiLong Fan, Jin Li, Xin Li, YuanHong Zhong, and JunWei Cao. Applying deep neural networks to the detection and space parameter estimation of compact binary coalescence with a network of gravitational wave detectors. *Science China Physics, Mechanics & Astronomy*, 62(6), February 2019.

[233] Chetan Verma, Amit Reza, Dilip Krishnaswamy, Sarah Caudill, and Gurudatt Gaur. Employing deep learning for detection of gravitational waves from compact binary coalescences. In *The 6th International Conference on Science and Technology (ICST21): Challenges and Opportunities for Innovation Research on Science Materials, and Technology in the Covid-19 Era*. AIP Publishing, 2022.

[234] A. Menéndez-Vázquez, M. Kolstein, M. Martínez, and Ll. M. Mir. Searches for compact binary coalescence events using neural networks in the ligo/virgo second observation period. *Phys. Rev. D*, 103:062004, Mar 2021.

[235] Shreejit Jadhav, Mihir Shrivastava, and Sanjit Mitra. Towards a robust and reliable deep learning approach for detection of compact binary mergers in gravitational wave data. *Machine Learning: Science and Technology*, 4(4):045028, nov 2023.

[236] Richard Qiu, Plamen G. Krastev, Kiranjyot Gill, and Edo Berger. Deep learning detection and classification of gravitational waves from neutron star-black hole mergers. *Physics Letters B*, 840:137850, 2023.

[237] Chayan Chatterjee, Linqing Wen, Foivos Diakogiannis, and Kevin Vinsen. Extraction of binary black hole gravitational wave signals from detector data using deep learning. *Phys. Rev. D*, 104:064046, Sep 2021.

[238] Vincent Boudart. Convolutional neural network to distinguish glitches from minute-long gravitational wave transients. *Phys. Rev. D*, 107:024007, Jan 2023.

[239] Naoki Koyama, Yusuke Sakai, Seiya Sasaoka, Diego Dominguez, Kentaro Somiya, Yuto Omae, Yoshikazu Terada, Marco Meyer-Conde, and Hirotaka Takahashi. Enhancing the rationale of convolutional neural networks for glitch classification in gravitational wave detectors: a visual explanation. *Mach. Learn. Sci. Tech.*, 5(3):035028, 2024.

[240] Daniel George, Hongyu Shen, and E. A. Huerta. Classification and unsupervised clustering of ligo data with deep transfer learning. *Phys. Rev. D*, 97:101501, May 2018.

[241] Patricia Schmidt, Mark Hannam, and Sascha Husa. Towards models of gravitational waveforms from generic binaries: A simple approximate mapping between precessing and nonprecessing inspiral signals. *Phys. Rev. D*, 86:104063, Nov 2012.

[242] The Event Horizon Telescope Collaboration. First sagittarius a* event horizon telescope results. vii. polarization of the ring. *The Astrophysical Journal Letters*, 964(2):L25, mar 2024.

[243] The Event Horizon Telescope Collaboration. First sagittarius a* event horizon telescope results. viii. physical interpretation of the polarized ring. *The Astrophysical Journal Letters*, 964(2):L26, mar 2024.

[244] William A. Fowler and F. Hoyle. Neutrino Processes and Pair Formation in Massive Stars and Supernovae. *Astrophys. J. Suppl.*, 9:201, December 1964.

[245] Z. Barkat, G. Rakavy, and N. Sack. Dynamics of supernova explosion resulting from pair formation. *Phys. Rev. Lett.*, 18:379–381, Mar 1967.

[246] J. R. Bond, W. D. Arnett, and B. J. Carr. The evolution and fate of Very Massive Objects. *The Astrophysics Journal*, 280:825–847, May 1984.

[247] S. E. Woosley, S. Blinnikov, and Alexander Heger. Pulsational pair instability as an explanation for the most luminous supernovae. *Nature*, 450(7168):390–392, November 2007.

[248] S. E. Woosley and Alexander Heger. The Pair-instability Mass Gap for Black Holes. *Astrophys. J. Lett.*, 912(2):L31, May 2021.

[249] Jenny E. Greene. Low-mass black holes as the remnants of primordial black hole formation. *Nature Communications*, 3(1), December 2012.

[250] Toshihiro Kawaguchi, Masahiro Kawasaki, Tsutomu Takayama, Masahide Yamaguchi, and Jun'ichi Yokoyama. Formation of intermediate-mass black holes as primordial black holes in the inflationary cosmology with running spectral index. *Monthly Notices of the Royal Astronomical Society*, 388(3):1426–1432, 07 2008.

[251] J. Aasi, B.P. Abbott, R. Abbott, T. Abbott, M.R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R.X. Adhikari, C. Affeldt, M. Agathos, N. Aggarwal, O.D. Aguiar, A. Ain, P. Ajith, A. Alemic, and B. Allen. Search for gravitational wave ringdowns from perturbed intermediate mass black holes in ligo-virgo data from 2005–2010. *Physical Review D*, 89(10), may 2014.

[252] LIGO Scientific Collaboration and Virgo Collaboration. Characterization of transient noise in advanced ligo relevant to gravitational wave signal gw150914. *Classical and Quantum Gravity*, 33(13):134001, jun 2016.

[253] Hunter Gabbard, Michael Williams, Fergus Hayes, and Chris Messenger. Matching matched filtering with deep networks for gravitational-wave astronomy. *Phys. Rev. Lett.*, 120:141103, Apr 2018.

[254] CunLiang Ma, Sen Wang, Wei Wang, and Zhoujian Cao. Using deep learning to predict matched signal-to-noise ratio of gravitational waves. *Phys. Rev. D*, 109:043009, Feb 2024.

[255] Hua-Mei Luo, Wenbin Lin, Zu-Cheng Chen, and Qing-Guo Huang. Extraction of gravitational wave signals with optimized convolutional neural network. *Frontiers of Physics*, 15(1):14601, February 2020.

[256] Vasileios Skliris, Michael R. K. Norman, and Patrick J. Sutton. Real-time detection of unmodelled gravitational-wave transients using convolutional neural networks, 2024. *arXiv e-Print Archive*, arXiv:2009.14611 [astro-ph.IM].

[257] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, page 41–48, New York, NY, USA, 2009. Association for Computing Machinery.

[258] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.

[259] Alex Graves, Abdel rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks, 2013. *arXiv e-Print Archive*, arXiv:1303.5778 [cs.NE].

[260] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, T. Akutsu, and G. Allen. Prospects for observing and localizing gravitational-wave transients with advanced ligo, advanced virgo and kagra. *Living Reviews in Relativity*, 23(1), September 2020.

[261] L. V. Kantorovich. Mathematical methods of organizing and planning production. *Manage. Sci.*, 6(4):366–422, July 1960.

[262] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.

[263] S. Kullback and R. A. Leibler. On information and sufficiency. *Ann. Math. Statist.*, 22(1):79–86, 1951.

[264] Derek Davis, Laurel V White, and Peter R Saulson. Utilizing aligo glitch classifications to validate gravitational-wave candidates. *Classical and Quantum Gravity*, 37(14):145001, June 2020.

[265] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'14, page 3320–3328, Cambridge, MA, USA, 2014. MIT Press.

[266] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. Convergent learning: Do different neural networks learn the same representations? In Dmitry Storcheus, Afshin Rostamizadeh, and Sanjiv Kumar, editors, *Proceedings of the 1st International Workshop on Feature Extraction: Modern Questions and Challenges at NIPS 2015*, volume 44 of *Proceedings of Machine Learning Research*, pages 196–212, Montreal, Canada, 11 Dec 2015. PMLR.

[267] Yi Pan, Alessandra Buonanno, Andrea Taracchini, Lawrence E. Kidder, Abdul H. Mroué, Harald P. Pfeiffer, Mark A. Scheel, and Béla Szilágyi.

Inspiral-merger-ringdown waveforms of spinning, precessing black-hole binaries in the effective-one-body formalism. *Physical Review D*, 89(8), April 2014.

[268] Héctor Estellés, Marta Colleoni, Cecilio García-Quirós, Sascha Husa, David Keitel, Maite Mateu-Lucena, Maria de Lluc Planas, and Antoni Ramos-Buades. New twists in compact binary waveform modeling: A fast time-domain model for precession. *Physical Review D*, 105(8), April 2022.

[269] Sascha Husa, Sebastian Khan, Mark Hannam, Michael Pürrer, Frank Ohme, Xisco Jiménez Forteza, and Alejandro Bohé. Frequency-domain gravitational waves from nonprecessing black-hole binaries. i. new numerical waveforms and anatomy of the signal. *Phys. Rev. D*, 93:044006, Feb 2016.

[270] Emmanuel A. Tassone, Paula A. Mandrilli, Carlos N. Kozameh, Gonzalo D. Quiroga, and José I. Nieva. Numerical evolution of the center of mass and angular momentum for binary black holes. *Phys. Rev. D*, 104:084038, Oct 2021.

[271] Stanislav Babak, Andrea Taracchini, and Alessandra Buonanno. Validating the effective-one-body model of spinning, precessing binary black holes against numerical relativity. *Physical Review D*, 95(2), January 2017.

[272] I. W. Harry, B. Allen, and B. S. Sathyaprakash. Stochastic template placement algorithm for gravitational wave data analysis. *Phys. Rev. D*, 80:104014, Nov 2009.

[273] P. Ajith, N. Fotopoulos, S. Privitera, A. Neunzert, N. Mazumder, and A. J. Weinstein. Effectual template bank for the detection of gravitational waves from inspiralling compact binaries with generic spins. *Phys. Rev. D*, 89:084041, Apr 2014.

[274] Ian Harry, Stephen Privitera, Alejandro Bohé, and Alessandra Buonanno. Searching for gravitational waves from compact binaries with precessing spins. *Phys. Rev. D*, 94:024012, Jul 2016.

[275] Connor McIsaac, Charlie Hoy, and Ian Harry. A search technique to observe precessing compact binary mergers in the advanced detector era, 2023. *arXiv e-Print Archive*, arXiv:2303.17364.

[276] M Cavaglià, S Gaudio, T Hansen, K Staats, M Szczepańczyk, and M Zanolin. Improving the background of gravitational-wave searches for core collapse supernovae: a machine learning approach. *Machine Learning: Science and Technology*, 1(1):015005, feb 2020.

[277] Teerth Gill Marco Cavaglia, Kai Staats. Finding the origin of noise transients in ligo data with machine learning. *Communications in Computational Physics*, 25(4):963–987, 2018.

[278] Serena Vinciguerra, John Veitch, and Ilya Mandel. Accelerating gravitational wave parameter estimation with multi-band template interpolation. *Class. Quant. Grav.*, 34(11):115006, 2017.

[279] Soichiro Morisaki. Accelerating parameter estimation of gravitational waves from compact binary coalescence using adaptive frequency resolutions. *Phys. Rev. D*, 104(4):044062, 2021.

[280] Rory Smith, Scott E. Field, Kent Blackburn, Carl-Johan Haster, Michael Pürrer, Vivien Raymond, and Patricia Schmidt. Fast and accurate inference on gravitational waves from precessing compact binaries. *Phys. Rev. D*, 94:044031, Aug 2016.

[281] Priscilla Canizares, Scott E. Field, Jonathan Gair, Vivien Raymond, Rory Smith, and Manuel Tiglio. Accelerated gravitational wave parameter estimation with reduced order modeling. *Phys. Rev. Lett.*, 114:071104, Feb 2015.

[282] Soichiro Morisaki, Rory Smith, Leo Tsukada, Surabhi Sachdev, Simon Stevenson, Colm Talbot, and Aaron Zimmerman. Rapid localization and inference on compact binary coalescences with the advanced ligo-virgo-kagra gravitational-wave detector network. *Phys. Rev. D*, 108:123040, Dec 2023.

[283] Lalit Pathak, Amit Reza, and Anand S. Sengupta. Fast and faithful interpolation of numerical relativity surrogate waveforms using meshfree approximation, 2024. *arXiv e-Print Archive*, arXiv:2403.19162 [gr-qc].

[284] Tanja Hinderer, Benjamin D. Lackey, Ryan N. Lang, and Jocelyn S. Read. Tidal deformability of neutron stars with realistic equations of state and their gravitational wave signatures in binary inspiral. *Phys. Rev. D*, 81:123016, Jun 2010.

[285] LIGO Scientific Collaboration and Virgo Collaboration. Gw190412: Observation of a binary-black-hole coalescence with asymmetric masses. *Phys. Rev. D*, 102:043015, Aug 2020.

[286] Geraint Pratten, Patricia Schmidt, Riccardo Buscicchio, and Lucy M. Thomas. Measuring precession in asymmetric compact binaries. *Phys. Rev. Res.*, 2:043096, Oct 2020.

[287] Charlie Hoy, Cameron Mills, and Stephen Fairhurst. Evidence for subdominant multipole moments and precession in merging black-hole-binaries from gwtc-2.1. *Phys. Rev. D*, 106:023019, Jul 2022.

[288] Tobias Blickle and Lothar Thiele. A comparison of selection schemes used in evolutionary algorithms. *Evolutionary Computation*, 4(4):361–394, 1996.

[289] F.P. Preparata and M.I. Shamos. Computational geometry: An introduction, 2012.

[290] Robert J. Fowler, Michael S. Paterson, and Steven L. Tanimoto. Optimal packing and covering in the plane are np-complete. *Information Processing Letters*, 12(3):133–137, 1981.

[291] Eric Thrane and Colm Talbot. An introduction to bayesian inference in gravitational-wave astronomy: Parameter estimation, model selection, and hierarchical models. *Publications of the Astronomical Society of Australia*, 36:e010, 2019.