


Universidade do Minho
Escola de Ciências

Oswaldo Gramaxo de Freitas **Gravitational Waves from Compact Binary Mergers**

Oswaldo Gramaxo de Freitas

**Gravitational Waves from
Compact Binary Mergers**



Universidade do Minho
Escola de Ciências

Oswaldo Gramaxo de Freitas

**Gravitational Waves from
Compact Binary Mergers**

Dissertação de Mestrado
Mestrado em Física

Trabalho efetuado sob a orientação do

Professor Doutor
António Joaquim Onofre Abreu Ribeiro Gonçalves


e do

Professor Doutor
José Antonio Font Roda

Direitos de autor e condições de utilização do trabalho por terceiros

Este é um trabalho académico que pode ser utilizado por terceiros desde que respeitadas as regras e boas práticas internacionalmente aceites, no que concerne aos direitos de autor e direitos conexos.

Assim, o presente trabalho pode ser utilizado nos termos previstos na licença abaixo indicada. Caso o utilizador necessite de permissão para poder fazer um uso do trabalho em condições não previstas no licenciamento indicado, deverá contactar o autor, através do RepositóriUM da Universidade do Minho.



<https://creativecommons.org/licenses/by/4.0/>

Acknowledgements

I must start by acknowledging the contribution of my advisors: Professor António Onofre, whose boundless enthusiasm got us into the field of gravitational wave (GW) astronomy in the first place and has been the kernel for many interesting discussions, and Professor Toni Font, whose expertise proved invaluable in introducing us to the field and its community. Both worked very hard to ensure I would have the opportunity to develop and disseminate relevant work, as well as providing chances to interact with the larger gravitational physics community. Another huge note of gratitude must go to Felipe Freitas, for guiding my first steps into machine learning and deep learning techniques. I want to also acknowledge the contributions of Alejandro Torres-Forné, António Morais, Nicolas Sanchis-Gual, Solange Nunes and João Álvares for various insights and contributions through our weekly group meetings. A word of appreciation goes to ~~minion~~ undergrad student Inês Rolo for contributing to the proofreading of chapter 4 in this dissertation.

On a more personal level I want to thank Filipa, Igor, Maura, Pedro, Beatriz, Cecília, Tito, and João for keeping me company in some way or other throughout a very inconvenient and pandemic (such a rude virus). Lastly, I wanted to give credit to the many great teachers I've had the luck of having from elementary school (where my teacher called me her little astrophysicist, funnily enough) all the way up to secondary school. If I have developed a reverence to curiosity and the pursuit of knowledge it is due to them, who always believed in me and pushed me to do my best. I am eternally grateful.

Statement of integrity

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the University of Minho.

Ondas Gravitacionais de Fusões Binárias Compactas

Resumo

As ondas gravitacionais (OG) são ondas no tecido do espaço-tempo causadas por fenómenos astrofísicos como supernovas de colapso do núcleo e fusões binárias compactas. Apesar das enormes energias envolvidas nestes eventos, a amplitude das OG geradas é notavelmente pequena, e a sua detecção directa não foi possível até muito recentemente: em 2015, os dois interferómetros da colaboração LIGO detectaram um evento astrofísico consistente com a fusão de dois buracos negros, dando início à era da astronomia de OG. A esta rede de detectores juntaram-se o interferómetro Virgo (Observatório Gravitacional Europeu) em 2017 e o interferómetro KAGRA (Universidade de Tóquio) em 2020, reforçando a nossa capacidade de detectar e analisar estes eventos.

A sondagem de regimes de gravidade extrema promete abrir novas portas para testes de física fundamental. Por exemplo, OG de fusões de estrelas de neutrões e supernovas permitem a exploração de processos nucleares no núcleo de estrelas compactas. Outras aplicações incluem a utilização de OG de sistemas compactos para testes de gravidade de campo forte e teorias de gravidade modificada. Futuros interferómetros espaciais, como LISA, poderão permitir um estudo da física do início do universo para além do que é possível com a radiação electromagnética. No entanto, as actuais abordagens à análise de dados baseiam-se em métodos de *match-filtering*, que têm custos computacionais significativos. As técnicas de *deep learning*, que têm tido uma explosão no desenvolvimento ao longo das últimas décadas, têm demonstrado a capacidade de obter elevados níveis de performance tanto em precisão como em tempo de execução. Estes métodos foram naturalmente propostos como um potencial aliado para a detecção e inferência de parâmetros nos sinais de OG.

Neste trabalho, exploramos a aplicação de redes neurais convolutivas em representações de tempo-frequência de OG de fusões binárias compactas, para a detecção de sinais e inferência de parâmetros. Mostramos que a detecção destes sinais pode ser alcançada com uma precisão significativa, e que a inferência de parâmetros físicos é viável nesta implementação.

Palavras-chave: ondas gravitacionais, binários compactos, deep learning, redes neuronais convolucionais, inferência de parâmetros

Gravitational Waves from Compact Binary Mergers

Abstract

Gravitational waves (GW) are ripples in the fabric of spacetime caused by astrophysical phenomena such as core-collapse supernovae and compact binary mergers. Despite the enormous energies involved in these events, the amplitude of the generated GWs is remarkably small, and their direct detection has been impossible until very recently: in 2015, the two interferometers of the LIGO collaboration detected an astrophysical event consistent with the merger of two black holes, ushering in a new era of GW astronomy. The detector network was joined by the European Gravitational Observatory’s Virgo interferometer in 2017 and the University of Tokyo’s KAGRA interferometer in 2020, further bolstering our capacity to detect and analyse these events.

Probing extreme gravity regimes promises to open new doors for tests of fundamental physics. For example, GWs from neutron star mergers and supernovae allow for the exploration of nuclear processes at the core of compact stars. Other applications include using GWs from compact systems for tests of strong-field gravity and modified gravity theories. Future space-based interferometers such as LISA may allow for a probing of early-universe physics beyond what is possible with electromagnetic radiation. However, current approaches to data analysis are based on match-filtering methods, which come with significant computational costs. Deep learning techniques, which have seen an explosion in development over the last couple of decades, have shown the capacity for high performance levels in both accuracy and speed at runtime. These methods have naturally been proposed as a potential ally for detection and parameter inference on GW signals.

In this work, we explore the application of convolutional neural networks on time-frequency representations of GW data from compact binary mergers, for both signal detection and parameter inference. We show that the detection of these signals can be achieved with significant accuracy, and that physical parameter inference is feasible in this implementation.

Keywords: gravitational waves, compact binaries, deep learning, convolutional neural networks, parameter inference

Contents

1	Introduction	15
2	Introduction to Gravitational Waves	19
2.1	The Einstein Equation	19
2.2	Gravitational waves in the linear approximation	20
2.2.1	Interactions with test masses in the TT gauge	22
2.3	Gravitational wave generation	23
2.3.1	Gravitational Waves generated by a binary system	27
2.4	Post Newtonian expansion and higher order contributions	33
3	Experimental setup	35
3.1	Michelson Interferometers	35
3.1.1	Interaction of a gravitational wave with a Michelson interferometer	37
3.1.2	Geometrical dependence of a GW interferometer detection	43
3.2	Noise in the advanced LIGO and advanced Virgo detectors	45
3.2.1	Seismic noise	45
3.2.2	Gravity gradients	46
3.2.3	Quantum noise	46
3.2.4	Suspension thermal noise	47
3.2.5	Coating brownian noise	47
3.2.6	Glitches	48
4	Machine Learning	49
4.1	Deep learning and neural networks	50
4.1.1	Convolutional Neural Networks	52
4.1.2	Residual Networks	54
4.2	Regularization	56
4.2.1	Weight decay	56
4.2.2	Dropout	57
5	Towards Bayesian Neural Networks	58
5.1	Bayesian modelling and the Gaussian Process	58
5.2	Variational Inference and MC dropout	60

6	Methods	62
6.1	Generating datasets	63
6.1.1	Single detector waveforms for classification	63
6.1.2	Multiple detector waveforms for classification	64
6.1.3	Mass dependence dataset	66
6.1.4	Regression datasets	66
6.2	Building a network for classification	67
6.2.1	Single-detector training distance comparison	67
6.2.2	Comparing single-detector and multiple-detector performance	69
6.3	Building a network for parameter estimation	71
6.3.1	Parameter estimation tests on synthetic injections	72
7	Results and discussion	80
7.1	Classification network	80
7.2	Regression network	82
7.2.1	GWTC-1	82
7.2.2	GWTC-2	84
8	Conclusions	87
	References	89

List of Figures

2.1	Evolution of the angular frequency over time for a coalescing binary system.	31
2.2	Gravitational waveform for a coalescing binary system.	33
3.1	Layout for a basic Michelson interferometer.	36
3.2	Layout of a Michelson interferometer with Fabry-Perot cavities.	42
3.3	Example of the geometry for GW detection.	44
3.4	Example map of the network antenna power pattern.	45
3.5	Design sensitivity curves for Advanced LIGO and Advanced Virgo.	46
3.6	Seismic isolation systems in gravitational wave interferometers.	47
4.1	Mock model of an FCNN.	51
4.2	Mock example of a 2-parameter loss function and its gradient.	52
4.3	Mock model of a CNN.	53
4.4	Example of a residual block.	55
4.5	Illustration of overfitting.	55
6.1	Example of a spectrogram produced for a simulation of the coalescence of a binary system.	62
6.2	Combining single detector spectrogram data into a single RGB image.	65
6.3	Performance of the different classifier networks.	69
6.4	Dependence of single detector signal scores in the component masses and luminosity distance.	70
6.5	Dependence of multiple detector signal scores in the component masses and luminosity distance.	70
6.6	ROC curve for the best-performing classifier.	71
6.7	Calibration results for $d_{\{L\}}$ using the different approximant datasets.	73
6.8	Calibration results for the network antenna power.	74
6.9	Calibration results for the chirp mass.	75
6.10	Calibration results for the effective inspiral spin.	76
6.11	Mean absolute error of distance predictions as a function of SNR.	77
6.12	Mean absolute error of network antenna power predictions as a function of SNR.	78
6.13	Mean absolute error of chirp mass predictions as a function of SNR.	78
6.14	Mean absolute error of effective inspiral spin predictions as a function of SNR.	79
7.1	Regression results for GWTC-1.	83
7.2	Regression results for GWTC-2.	85

List of Tables

6.1	Description of the classification datasets.	67
6.2	Description of the regression datasets.	68
6.3	Convolutional Neural Networks architectures employed for classification and regression tasks.	72
7.1	Classifier scores for GWTC-1 and GWTC-2.	81

Nomenclature

AI	Artificial Intelligence
ASD	Amplitude spectral density
BBH	Binary black hole
BC	Before Christ
CBC	Compact binary coalescence
CNN	Convolutional neural network
CV	Computer Vision
DL	Deep learning
EFE	Einstein field equations
FCNN	Fully connected neural networks
FP	Fabry-Perot
GP	Gaussian process
GPS	Global positioning system
GR	General relativity
GW	Gravitational waves
IBM	International Business Machines
ISCO	Innermost stable circular orbit
KL	Kullback-Leibler
LIGO	Laser Interferometer Gravitational-Wave Observatory
LVC	LIGO/Virgo collaboration
MC	Monte Carlo
ML	Machine learning

MSE	Mean squared error
NLP	Natural language processing
NN	Neural network
PN	Post-Newtonian
ResNet	Residual network
RGB	Red, Green, Blue (referring to the color channels in a digital image)
ROC	Receiver operating characteristic
USA	United states of America
VI	Variational inference

1 Introduction

“It’s a warm summer evening, circa 600 BC. You’ve finished your shopping at the local market, or agora... and you look up at the night sky. There you notice some of the stars seem to move, so you name them *planetes* or wanderer.” - Sheldon Cooper

The observation of the skies has enchanted the human spirit since the depths of prehistory. From star charts on ornate disks [1] to astronomically aligned megalithic structures of white quartz cobblestone [2], mankind’s fascination and veneration of celestial objects and their behaviour is a leitmotif through time.

Our first hint at the usage of mathematics applied to astronomical events comes from the banks of the Euphrates, where Babylonian scholars studied the length of days and nights over the year [3]. Indeed the Babylonians gave us also the first exploration of the periodicity of planetary movement [4]. However, they were not solitary in their endeavours, and many ancient societies would go on to formalize their studies of the heavens. Of these, the most influential to the development of western astronomy were the ancient Greeks, who placed astronomy among the four mathematical arts (along with arithmetic, geometry, and music). The works of many Greek (and later, Roman) natural philosophers and mathematicians, such as Claudius Ptolemaeus’s *Almagest* (whose geocentric system superceded earlier efforts by Aristarchus of Samos and Seleucus of Seleucia to propose a heliocentric system), found echoes throughout the Christian and Islamic worlds well into the late medieval period [5, 6]. It was during the Renaissance period, however, that astronomy met physics, as Isaac Newton’s *Principia* revolutionised the mathematical treatment of physical phenomena, managing to accurately explain most astronomical predictions of the time under the solid theoretical framework of classical mechanics, the inverse-square law of universal gravitation being a particular highlight. Newton’s theory was and is immensely influential, and is still widely used in astronomy for many calculations [7]. However, this theory was still not perfect: observations of the orbit of mercury showed an acceleration of the perihelion that could not be described in the theory [8].

Proposed by Albert Einstein in 1915, the theory of general relativity (GR) would come to be known for its exquisite mathematical elegance and predictive power. Positing gravitational interactions as a deformation of an underlying spacetime structure, rather than mediated by a force such as the one in Newtonian gravity, GR’s geometrical nature, sense of scale, and initially unintuitive coupling of space and time have not only greatly advanced our knowledge of the cosmos, but have also stoked the flames of popular imagination. GR

was given mainstream credibility by Einstein’s general relativistic 1915 description of the anomalous perihelion shift of Mercury, showing a predictive power beyond Newtonian gravity, and, most famously, by the Eddington experiment’s confirmation of GR’s predicted values for gravitational lensing (twice the value of previous Newtonian predictions) [9]. A plethora of further observations have reinforced its place as a remarkably precise tool and a centerpiece of modern physics [10].

One of the most exotic phenomena predicted by GR is the existence of black holes, regions of spacetime wherein no causal paths can be used to escape. Similar objects had been alluded to in the Newtonian framework: in 1784, John Michell proposed the idea of “dark stars”, stars whose escape velocity exceeded the speed of light for objects with sufficient mass and density [11], an idea which was also independently proposed by the mathematician Pierre-Simon Laplace in 1796, with a proof being made available in 1799[12, 13]. Indeed, in 1916, Schwarzschild’s solution to the Einstein field equations (EFE) for a spherically symmetric, non-rotating body pointed to a singularity in the metric for a radius $r_s = \frac{2MG}{c^2}$ (the Schwarzschild radius), depending on a choice of integration constant, where the curvature of spacetime would shoot to infinity [14]. However, a revised version of this solution by Hilbert, using a different coordinate system, also predicted a singularity at the center of the region, at $r = 0$ [15, 16]. These singularities were present in subsequent spherical solutions to the EFE, and as such were widely studied with the eventual realization that the singularity at r_s was merely an artifact of the coordinate system used [17]. The singularity at the center remained nonetheless unsurmountable, until Stephen Hawking and Roger Penrose developed the Hawking-Penrose singularity theorems *circa* 1965, which proved that this point of infinite curvature was a robust prediction in General Relativity [18, 19]. Objects under these conditions fit the modern definition of a black hole, and Penrose’s work would come to be honored with half the 2020 Nobel Prize in Physics [20].

Another exciting prediction of GR is the existence of gravitational waves (GW): ripples in the fabric of space and time resulting from the dynamics of massive systems or bodies that meet certain requirements. The gravitational field of an object has no charge, so waves cannot be produced by dipole objects like in electromagnetism. However, gravitational systems possessing *quadrupole* momentum radiate gravitational waves. Significant sources include compact binary systems and supernovae.

In 1916, Einstein made the initial proposal for the existence of three types of GWs [21]; however, in 1922, Eddington showed two of those types of waves to be simply an artifact of the coordinate system used, and indeed that they could move at arbitrary speed. The third type (transverse-transverse), however, was not removable through such means, and in

fact was shown to travel at the speed of light [22]. Development of the subject was left without major breakthroughs for a number of decades, during which Einstein himself cast some doubt over the existence of this third wave [23]. In 1956 Felix Pirani finally showed the existence of these waves in a coordinate-independent manner [24]. This, combined with Richard Feynman’s “sticky bead argument”, which showed that a passing GW could indeed transfer energy to a system, shifted the general scientific consensus towards the acceptance of GWs as meaningful predictions of general relativity [25]. Despite some initial attempts by Weber, in the 70s, to directly detect these waves, the first experimental hint of their existence was obtained indirectly: in 1975, Hulse and Taylor published a paper announcing the discovery of the first binary pulsar, where they realized that the periodicity of radio emissions from the pulsar they were studying meant that it was in fact orbiting around the center of mass of a binary system, in which the second component was a neutron star [26]. Then, a 1981 paper by Weisberg, Taylor and Fowler found that the orbit of the binary was contracting in a way that was remarkably consistent with the loss of energy from such a system through GWs [27]. These works amplified the efforts to develop facilities suitable for direct detections, and ended up earning Hulse and Taylor the Nobel prize for physics in 1993 [28].

The idea of using a Michelson-style interferometer with hanging mirrors dates back to 1962, when Michael Gertsenshtein and Vladislav Pustovoit, in Moscow, proposed just such an idea. This was also independently proposed by Weber and Weiss a few years later [29]. In 1994, the LIGO project broke ground in the USA, while the Virgo interferometer in Europe started construction in Cascina, Italy in 1996. Their respective operations started in 2002 and 2007. The initial runs yielded no detections, which was not surprising given the sensitivity of the detectors at the time. The instruments were further improved, giving way to Advanced LIGO [30] and Advanced Virgo [31], and on the 14th of September of 2015, the first binary black hole (BBH) event, GW150914, was detected during an engineering run - the herald to the era of GW astronomy. The nature of GWs, given gravity’s weak coupling to other forms of energy, means that these signals should be relatively undisturbed in the path from their source to the observer, and as such GW astronomy can offer insights into objects which other tools, such as X-ray telescopes or cosmic ray detectors, are not able to due to cosmic interference. Furthermore, GW astronomy offers the enticing possibility of observing exotic objects which do not interact significantly through Standard Model forces, leaving only a gravitational signature, such as Proca stars [32].

This dissertation intends to explore novel ways of detecting and interpreting data from GW interferometers, specifically taking advantage of the exponential progress recently made

in the application of neural networks for image processing in order to explore 2D representations of GW data. The document will be split into six chapters: first, we will introduce concepts from general relativity and derive the expression for the GW strain in the linear approximation, as well as make some comments on alternate ways of obtaining such results. A quick discussion of Michelson-Morley laser interferometers is had, and we investigate the specific implementation of the concept in the LIGO and Virgo facilities. After that we will introduce some basic concepts from machine learning and neural networks, and consider implementations of such techniques. Then we shall explore the work performed in building a neural network that is capable of detecting the presence of GW signatures in spectrogram data. Subsequently we will present a similar construction to perform parameter estimation on GW events. Finally we will discuss results and consider possible future work.

2 Introduction to Gravitational Waves

Throughout this work, the signature used for the metric tensor is $(-+++)$. We will denote partial derivatives as $A_{\mu,\nu} = \partial_\nu A_\mu$ and covariant derivatives as $A_{\mu;\nu} = \nabla_\nu A_\mu$. Furthermore we will be using the D'Alembert operator $\square A \equiv \partial^\mu \partial_\mu A$. Note that Greek indices run from 0 to 3, while Latin indices run from 1 to 3.

2.1 The Einstein Equation

A first description of gravitational waves can be given by taking a linear approximation of GR far away from a GW source. In this section we will show how to obtain the Einstein equations and develop this approximation.

The gravitational action is given by the sum of a curvature term and a mass term:

$$S = S_E + S_M = \frac{1}{c} \int d^4x \left[\frac{c^4}{16\pi G} R + \mathcal{L}_M \right] \sqrt{-g}, \quad (2.1)$$

where the constants c and G stand for the speed of light and the Newtonian constant of gravitation, respectively. The symbol g , with no scripts, stands for the determinant of the metric tensor $g_{\mu\nu}$, and R stands for the Ricci scalar, which is given by the trace of the Ricci tensor $R_{\mu\nu}$. From the principle of least action we impose $\delta S = 0$, and so:

$$\begin{aligned} \delta S &= \frac{1}{c} \int d^4x \left[\frac{c^4}{16\pi G} \frac{\delta(R\sqrt{-g})}{\delta g^{\mu\nu}} + \frac{\delta(\mathcal{L}_M\sqrt{-g})}{\delta g^{\mu\nu}} \right] \\ &= \frac{1}{c} \int d^4x \left[\frac{c^4}{16\pi G} \left(\sqrt{-g} \frac{\delta R}{\delta g^{\mu\nu}} + R \frac{\delta \sqrt{-g}}{\delta g^{\mu\nu}} \right) + \frac{\delta(\mathcal{L}_M\sqrt{-g})}{\delta g^{\mu\nu}} \right] = 0. \end{aligned} \quad (2.2)$$

If we define

$$T_{\mu\nu} = -\frac{2}{\sqrt{-g}} \frac{\delta(\mathcal{L}_M\sqrt{-g})}{\delta g^{\mu\nu}}$$

as the stress-energy tensor, then

$$\int d^4x \frac{c^4}{16\pi G} \left(\sqrt{-g} \frac{\delta R}{\delta g^{\mu\nu}} + R \frac{\delta \sqrt{-g}}{\delta g^{\mu\nu}} \right) = \int d^4x \frac{\sqrt{-g}}{2} T_{\mu\nu},$$

where the variation of R is given by:

$$\delta R = \delta(g^{\mu\nu} R_{\mu\nu}) = R_{\mu\nu} \delta g^{\mu\nu} + (g^{\mu\nu} \delta \Gamma_{\mu\nu}^\rho + g^{\mu\rho} \delta \Gamma_{\nu\mu}^\nu)_{;\rho},$$

where we use the Christoffel symbols $\Gamma_{\mu\nu}^{\sigma} = (1/2) g^{\alpha\sigma} (g_{\alpha\mu,\nu} + g_{\alpha\nu,\mu} - g_{\mu\nu,\alpha})$. However, as δR is being integrated on an arbitrary 4-volume, the second term, which is a total derivative, generates only a boundary term and as such does not contribute to the dynamics. On the other hand, the variation of $\sqrt{-g}$ is given by

$$\delta\sqrt{-g} = -\frac{1}{2\sqrt{-g}}\delta g = -\frac{1}{2}\sqrt{-g}g_{\mu\nu}\delta g^{\mu\nu}.$$

We can therefore rewrite (2.2) as:

$$\int d^4x \frac{c^4}{16\pi G} \left(R_{\mu\nu}\sqrt{-g} - \frac{1}{2}R\sqrt{-g}g_{\mu\nu} \right) = \int d^4x \frac{\sqrt{-g}}{2} T_{\mu\nu},$$

and since the integration is made in an arbitrary 4-volume, we thus obtain the Einstein equations:

$$R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \quad (2.3)$$

2.2 Gravitational waves in the linear approximation

A linear approximation to the Einstein equations can be taken by formulating the metric as a small perturbation on curved spacetime,

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

It is important to note that this approximation does not have the same symmetries as the original equations. In particular, under the coordinate transformation $x^\mu \rightarrow x'^\mu = x^\mu + \xi^\mu(x)$, the metric transforms as:

$$h_{\mu\nu} \rightarrow h'_{\mu\nu} = h_{\mu\nu} - (\xi_{\nu,\mu} + \xi_{\mu,\nu}). \quad (2.4)$$

It is useful at this point to define $h = h^\mu_\mu$ and

$$\begin{aligned} \bar{h}_{\mu\nu} &= h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h, \\ h_{\mu\nu} &= \bar{h}_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\bar{h}. \end{aligned}$$

From the linearization of the Riemann tensor, $R_{\rho\sigma\mu\nu} = \partial_\mu\Gamma_{\rho\nu\sigma} - \partial_\nu\Gamma_{\rho\mu\sigma} + \Gamma_{\rho\mu\lambda}\Gamma_{\nu\sigma}^\lambda - \Gamma_{\rho\nu\lambda}\Gamma_{\mu\sigma}^\lambda$, we obtain

$$R_{\mu\nu\rho\sigma} = \frac{1}{2} [h_{\nu\rho,\mu\sigma} + h_{\mu\sigma,\nu\rho} - h_{\nu\sigma,\mu\rho} - h_{\mu\rho,\nu\sigma}] .$$

Plugging this into (2.3), we get

$$\square\bar{h}_{\mu\nu} + \eta_{\mu\nu}\bar{h}^{\rho\sigma}_{;\rho\sigma} - \bar{h}^{\rho}_{\mu\rho,\nu} - \bar{h}^{\rho}_{\nu\rho,\mu} = \frac{16\pi G}{c^4}T_{\mu\nu} . \quad (2.5)$$

Since the metric tensor is symmetric this results in ten independent equations to be solved. However, using what we know from (2.4) we can see that

$$\bar{h}_{\mu\nu} \rightarrow \bar{h}'_{\mu\nu} = \bar{h}_{\mu\nu} - (\xi_{\nu,\mu} + \xi_{\mu,\nu} - \eta_{\mu\nu}\xi^{\rho}_{;\rho}) , \quad (2.6)$$

$$\bar{h}^{\nu}_{\mu\nu} \rightarrow (\bar{h}^{\nu}_{\mu\nu})' = \bar{h}^{\nu}_{\mu\nu} - \square\xi_{\mu} , \quad (2.7)$$

and so we can choose an arbitrary ξ_{μ} to transform our coordinates such that $\square\xi_{\mu} = \bar{h}^{\nu}_{\mu\nu}$, or equivalently:

$$\bar{h}^{\nu}_{\mu\nu} = 0 . \quad (2.8)$$

This is known as the Lorentz gauge, and it gives the linearized Einstein equations the form of a simple wave equation:

$$\square\bar{h}_{\mu\nu} = \frac{16\pi G}{c^4}T_{\mu\nu} . \quad (2.9)$$

While in general we would have ten independent components in the metric (owing to $g_{\mu\nu}$ being symmetric), the imposition of the Lorentz gauge gives us 4 conditions, and therefore we can reduce the number of independent tensor components to six.

In regions away from the source of the gravitational perturbation, $T_{\mu\nu} = 0$, so we can write a simple solution to the equations,

$$\square\bar{h}_{\mu\nu} = 0 \iff \left[-\frac{1}{c^2}\partial_t^2 + \nabla^2 \right] \bar{h}_{\mu\nu} = 0 ,$$

which immediately highlights that these gravitational waves travel at the speed of light. However, we still have some gauge freedom in our problem that we can take advantage of. Condition (2.8) still allows a further coordinate transformation $x^\mu \rightarrow x'^\mu = x^\mu + \xi^\mu(x)$ so long as $\square\xi_{\mu} = 0$, as according to equation (2.7) this will not alter the Lorentz gauge condition. $\square\xi_{\mu} = 0$ also means $\square(\xi_{\nu,\mu} + \xi_{\mu,\nu} - \eta_{\mu\nu}\xi^{\rho}_{;\rho}) = 0$, so we can see from equation (2.6) and equation (2.9) that this second gauge choice will not alter the form of our equations.

Now we can apply this coordinate transformation once for each coordinate x^μ , which

again according to equation (2.6) allows us to impose another four conditions on $\bar{h}_{\mu\nu}$, thus reducing the number of independent components of the tensor to two. A convenient choice is to pick ξ^0 such that the trace of \bar{h} vanishes, in which case we have $\bar{h}_{\mu\nu} = h_{\mu\nu}$. We can then choose our three remaining functions ξ^i such that $h^{0i} = 0$. An important consequence of this is that the Lorentz condition for $\mu = 0$ gives us $\bar{h}_{0\nu}{}^\nu = -(1/c)\partial^t h_{00} + \partial^i h_{0i} = 0$, but since $h_{0i} = 0$ we are simply left with

$$\partial^t h_{00} = 0,$$

which tells us that under these choices h_{00} becomes constant in time. This can be seen as the static part of gravitational interaction (the gravitational potential of the source). Since we are concerned with the dynamic aspect that gives rise to GWs, for the treatment of this problem we can ignore h_{00} by setting it to zero.

The conditions we have set on our problem,

$$h^{0\mu} = 0, \quad h_i^i = 0, \quad \partial^j h_{ij} = 0,$$

define the transverse-traceless (or TT) gauge, wherein we will refer to the metric perturbation as $h_{\mu\nu}^{TT}$.

2.2.1 Interactions with test masses in the TT gauge

In this gauge, equation (2.9) yields plane wave solutions of the type $h_{ij}^{TT}(x) = e_{ij}(\mathbf{k})e^{ik^\mu x_\mu}$, where $k^\mu = (\omega/c, \mathbf{k})$, $\omega/c = |\mathbf{k}|$ and e_{ij} is the polarization tensor (we have two polarizations, one for each degree of freedom). Since for a plane wave the condition $\partial^j h_{ij} = 0$ is equivalent to $n^j h_{ij} = 0$ with $n^j = k^j/\mathbf{k}$, we can choose coordinates such that n is along z , and so we can write h^{TT} as

$$h_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}_{\mu\nu} \cos \left[\omega \left(t - \frac{z}{c} \right) \right],$$

or, simplifying with regard to the number of degrees of freedom,

$$h_{ij}^{TT} = \begin{pmatrix} h_+ & h_\times & 0 \\ h_\times & -h_+ & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij} \cos \left[\omega \left(t - \frac{z}{c} \right) \right],$$

where h_+ and h_\times are known as the “plus” and “cross” polarizations due to their effect on rings of particles (which will not be explored here). Note that since in the TT gauge we have $h^{0\mu} = 0$, we can always refer to the tensor with latin indices to specify that we are treating only the spatial part. We can now define the spacetime interval for the metric $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}^{TT}$ as

$$ds^2 = -c^2 dt^2 + dz^2 + \left(1 + h_+ \cos\left[\omega\left(t - \frac{z}{c}\right)\right]\right) dx^2 + \left(1 - h_+ \cos\left[\omega\left(t - \frac{z}{c}\right)\right]\right) dy^2 + 2h_\times \cos\left[\omega\left(t - \frac{z}{c}\right)\right] dx dy.$$

The proper distance s between two test masses at $s_a = (t, x_a, 0, 0)$ and $s_b = (t, x_b, 0, 0)$ will be given by

$$\Delta s = \int_{s_a}^{s_b} ds = \int_{x_a}^{x_b} dx \sqrt{1 + h_+ \cos(\omega t)} = |x_b - x_a| \sqrt{1 + h_+ \cos(\omega t)}, \quad (2.10)$$

and so we can see that under the effect of gravitational waves there is a periodic oscillation to the proper distance between test masses. Since the travel time of light is given by $\Delta\tau = \frac{\Delta s}{c}$, these small movements will cause phase shifts that can be detected using Michelson-Morley style interferometers with sufficient sensitivity, as in the case of the LIGO/Virgo detectors which will be explored later on.

2.3 Gravitational wave generation

For a gravitationally bound two-body system with reduced mass μ and total mass M , the kinetic energy $K = -U/2$, where U is the gravitational potential. This translates to

$$\frac{1}{2}\mu v^2 = \frac{1}{2} \frac{G\mu M}{r} = \frac{\mu c^2 r_s}{4 r},$$

or, rearranging:

$$\frac{v^2}{c^2} = \frac{r_s}{2r}.$$

We take weak gravitational fields to imply $r \gg r_s$, and so also $v \ll c$. This means that for gravity-dominated systems the linear weak field approximation implies also a low-velocity approximation.

We now look again to the linearized EFE in the Lorentz gauge given in equation (2.9),

$$\square \bar{h}_{\mu\nu} = -\frac{16\pi G}{c^4} T_{\mu\nu}.$$

This linear equation can be solved using Green's functions, so that

$$\bar{h}_{\mu\nu}(x) = -\frac{16\pi G}{c^4} \int d^4x' G(x-x') T_{\mu\nu}(x').$$

As we are looking at a radiation problem, as in electromagnetism, the appropriate solution to use is the retarded Green's function

$$G(x-x') = -\frac{1}{4\pi |\mathbf{x}-\mathbf{x}'|} \delta\left(\left[t - \frac{|\mathbf{x}-\mathbf{x}'|}{c}\right] - t'\right),$$

wherein x' is limited to the source region. The solution to equation (2.9) is then

$$\bar{h}_{\mu\nu}(t, \mathbf{x}) = \frac{4G}{c^4} \int \frac{d^3\mathbf{x}'}{|\mathbf{x}-\mathbf{x}'|} T_{\mu\nu}\left(t - \frac{|\mathbf{x}-\mathbf{x}'|}{c}, \mathbf{x}'\right). \quad (2.11)$$

Now we would like to take a look at this solution away from the source, in the TT gauge. For that we will need to define a projection operator to apply to the obtained solution. We will first define the tensor $P_{ij}(\mathbf{n}) = \delta_{ij} - \mathbf{n}_i \mathbf{n}_j$, which is in itself a projector (note that $P_{ij} P_{jk} = P_{ik}$). Then we can construct the Lambda tensor

$$\Lambda_{ijkl} = P_{ik} P_{jl} - \frac{1}{2} P_{ij} P_{kl}, \quad (2.12)$$

which has the property which we seek, that is,

$$h_{ij}^{TT}(x) = \Lambda_{ijkl}(\mathbf{n}) h_{kl}(x).$$

Using the lambda tensor, and keeping in mind that $h_{\mu\nu}^{TT} = \bar{h}_{\mu\nu}^{TT}$ we can then obtain the solution to our problem in the TT gauge,

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{4G}{c^4} \Lambda_{ijkl}(\mathbf{n}) \int \frac{d^3\mathbf{x}'}{|\mathbf{x}-\mathbf{x}'|} T_{kl}\left(t - \frac{|\mathbf{x}-\mathbf{x}'|}{c}, \mathbf{x}'\right). \quad (2.13)$$

Defining $r = |\mathbf{x}|$ and d as the radius of the source, and knowing that far away from the

source $|\mathbf{x} - \mathbf{x}'| = r - \mathbf{x}' \cdot \mathbf{n} + O(d^2/r) \approx r$, we can expand T_{kl} around $(\mathbf{x}' \cdot \mathbf{n})/c$ as:

$$T_{kl} \left(t - \frac{r}{c} + \frac{\mathbf{x}' \cdot \mathbf{n}}{c}, \mathbf{x}' \right) \approx \left[T_{kl} + (\mathbf{x}' \cdot \mathbf{n}) \partial_0 T_{kl} + \frac{1}{2} (\mathbf{x}' \cdot \mathbf{n})^2 \partial_0^2 T_{kl} + \dots \right]_{x=(t-\frac{r}{c}, \mathbf{x})}$$

We can now define the momenta of T_{ij} :

$$\begin{aligned} S_{ij}(t) &= \int d^3x T_{ij}(t, \mathbf{x}), \\ S_{ij}{}^k(t) &= \int d^3x T_{ij}(t, \mathbf{x}) x^k, \\ S_{ij}{}^{kl}(t) &= \int d^3x T_{ij}(t, \mathbf{x}) x^k x^l, \\ S_{ij}{}^{kl\dots z}(t) &= \int d^3x T_{ij}(t, \mathbf{x}) x^k x^l \dots x^z. \end{aligned}$$

and with this, we can rewrite equation (2.13) as

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{4G}{rc^4} \Lambda_{ijkl}(\mathbf{n}) \left[S_{kl}(x) + n_m \partial_0 S_{kl}{}^m(x) + \frac{1}{2} n_m n_p \partial_0^2 S_{kl}{}^{mp}(x) + \dots \right]_{x=(t-\frac{r}{c}, \mathbf{x})}. \quad (2.14)$$

It is now common to define the momenta of the energy density:

$$\begin{aligned} M &= \frac{1}{c^2} S_{00}, \\ M^i &= \frac{1}{c^2} S_{00} x^i, \\ M^{ij} &= \frac{1}{c^2} S_{00} x^i x^j. \end{aligned}$$

Using the fact that, from conservation of momentum $\partial^\nu T_{\mu\nu} = 0 \iff \partial^0 T_{\mu 0} = -\partial^i T_{\mu i}$:

$$\begin{aligned} \dot{M}^{ij} &= c \partial_0 M^{ij} = \frac{1}{c} \int d^3x \partial_0 (T_{00}(t, \mathbf{x}) x^i x^j) \\ &= -\frac{1}{c} \int d^3x (\partial_k T_{0k}(t, \mathbf{x})) x^i x^j \\ &\updownarrow (\text{integration by parts}) \updownarrow \\ &= \frac{1}{c} \int d^3x T_{0k}(t, \mathbf{x}) \left(\underbrace{(\partial_k x^i)}_{\delta_k^i} x^j + x^i \underbrace{(\partial_k x^j)}_{\delta_k^j} \right) \end{aligned}$$

$$= \frac{1}{c} \int d^3x (T_{0i}(t, \mathbf{x})x^j + T_{0j}(t, \mathbf{x})x^i) ,$$

and so we can arrive at the result

$$\begin{aligned} \ddot{M}^{ij} &= \int d^3x \partial_0 (T_{0i}(t, \mathbf{x})x^j + T_{0j}(t, \mathbf{x})x^i) \\ &= \int d^3x (T_{ji}(t, \mathbf{x}) + T_{ij}(t, \mathbf{x})) . \end{aligned}$$

Since $T_{\mu\nu}$ is symmetric, we have

$$\begin{aligned} \ddot{M}^{ij} &= 2 \int d^3x T_{ij}(t, \mathbf{x}) , \\ \frac{1}{2} \ddot{M}^{ij} &= S_{ij} , \end{aligned}$$

which we can now substitute into the leading order of equation (2.14):

$$h_{ij}^{TT}(t, \mathbf{x}) = \frac{2G}{rc^4} \Lambda \ddot{M}_{ijkl}^{12}(\mathbf{n}) \ddot{M}^{kl}(t - \frac{r}{c}, \mathbf{x}) . \quad (2.15)$$

We now want to take care of the effect of the lambda tensor on \ddot{M}^{kl} . To that end, we first look back at the definition of the projector $P_{ij}(\mathbf{n}) = \delta_{ij} - \mathbf{n}_i \mathbf{n}_j$. If we take $\mathbf{n} = \hat{\mathbf{z}} = (0, 0, 1)$, then it is easy to see that there are no non-diagonal terms and that for P_{zz} the kronecker delta term and $n_z n_z = 1$ cancel out, and so

$$P_{ij} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij} .$$

We can then calculate $\Lambda_{ijkl}(\mathbf{n}) \ddot{M}^{kl} = [P_{ik} P_{jl} - \frac{1}{2} P_{ij} P_{kl}] \ddot{M}^{kl}$, where

$$P_{ik} P_{jl} \ddot{M}^{kl} = \begin{pmatrix} \ddot{M}^{11} & \ddot{M}^{12} & 0 \\ \ddot{M}^{21} & \ddot{M}^{22} & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij}$$

and

$$\frac{1}{2} P_{ij} P_{kl} \ddot{M}^{kl} = \frac{1}{2} \begin{pmatrix} \ddot{M}^{11} + \ddot{M}^{22} & 0 & 0 \\ 0 & \ddot{M}^{11} + \ddot{M}^{22} & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij} .$$

Plugging these terms into place, and taking into account that $\ddot{M}^{12} = \ddot{M}^{21}$, we then get:

$$\Lambda_{ijkl}(\hat{\mathbf{z}})\ddot{M}^{kl} = \begin{pmatrix} \frac{\ddot{M}^{11}-\ddot{M}^{22}}{2} & \ddot{M}^{12} & 0 \\ \ddot{M}^{12} & -\frac{\ddot{M}^{11}-\ddot{M}^{22}}{2} & 0 \\ 0 & 0 & 0 \end{pmatrix}_{ij},$$

which in tandem with (2.15) gives us the gravitational wave polarization amplitudes

$$h_+ = \frac{G}{rc^4} (\ddot{M}^{11} - \ddot{M}^{22}),$$

$$h_\times = \frac{2G}{rc^4} \ddot{M}^{12}.$$

2.3.1 Gravitational Waves generated by a binary system

We will now use the developed formalism to treat the problem of gravitational waves emitted by a general astrophysical binary system. It's important to note that at this point we are treating the bodies as point masses with no structure. This is a good enough approximation as long as the bodies are well enough separated. Another approximation we will be using for the sake of simplicity is the consideration of orbits with no eccentricity (i.e. circular orbits). We can feel comfortable about this as binary systems with eccentric orbits tend to decay to circular orbits due to the energy loss from GW emission [33, 34].

Consider then two bodies with masses m_1 and m_2 , at positions \mathbf{r}_1 and \mathbf{r}_2 . We define the separation $R = |\mathbf{r}_2 - \mathbf{r}_1|$ and assume a circular orbit with orbital frequency

$$\omega = \sqrt{G \frac{m_1 + m_2}{R^3}}. \quad (2.16)$$

If we choose our coordinates such that the orbit lies on the (x, y) plane, then in the center-mass frame we have:

$$\mathbf{r}_1 = \frac{\mu}{m_1} R (\cos(\omega t), \sin(\omega t), 0),$$

$$\mathbf{r}_2 = -\frac{\mu}{m_2} R (\cos(\omega t), \sin(\omega t), 0), \quad (2.17)$$

where $\mu = m_1 m_2 / m$ is the reduced mass of the system and $m = m_1 + m_2$.

Note that for the weak-field and non-relativistic regime we have $T_{00} = \rho c^2$ and equivalently $M = \int d^3\mathbf{r} \rho(\mathbf{r})$, where $\rho(\mathbf{r})$ is the mass density (outside the mentioned regime there will be further contributions to T_{00}). For the system described above, where the components are

point masses, we have

$$\rho(\mathbf{r}) = m_1 \delta(\mathbf{r} - \mathbf{r}_1) + m_2 \delta(\mathbf{r} - \mathbf{r}_2) .$$

Let us then calculate the mass quadrupole

$$\begin{aligned} M^{ij} &= \frac{1}{c^2} \int d^3\mathbf{r} T_{00}(t, \mathbf{r}) \mathbf{r}^i \mathbf{r}^j \\ &= \int d^3\mathbf{r} \rho(\mathbf{r}) \mathbf{r}^i \mathbf{r}^j \\ &= \int d^3\mathbf{r} [m_1 \delta(\mathbf{r} - \mathbf{r}_1) + m_2 \delta(\mathbf{r} - \mathbf{r}_2)] \mathbf{r}^i \mathbf{r}^j \\ &= m_1 (\mathbf{r}_1)_i (\mathbf{r}_1)_j + m_2 (\mathbf{r}_2)_i (\mathbf{r}_2)_j + C , \end{aligned}$$

such that

$$\begin{aligned} M^{11} &= \left(\frac{1}{m_1} + \frac{1}{m_2} \right) \mu^2 R^2 \cos^2(\omega t) + C \\ &= \mu R^2 \cos^2(\omega t) + C \\ &= \frac{\mu R^2}{2} [1 + \cos(2\omega t)] + C \\ &= \frac{\mu R^2}{2} \cos(2\omega t) + C , \\ M^{22} &= \mu R^2 \sin^2(\omega t) + C \\ &= \mu R^2 [1 - \cos^2(\omega t)] + C \\ &= -\frac{\mu R^2}{2} \cos(2\omega t) + C , \\ M^{12} &= \mu R^2 \cos(\omega t) \sin(\omega t) \\ &= \frac{\mu R^2}{2} \sin(2\omega t) = M^{21} . \end{aligned}$$

With this, we can now apply the second time derivative to these elements to ultimately obtain the polarization amplitudes. If we further use the relation $R = (Gm/\omega^2)^{1/3}$, and define the chirp mass $\mathcal{M} = \mu^{3/5} m^{2/5}$, then we have

$$\begin{aligned} h_+(t) &= \frac{G}{rc^4} \left(\ddot{M}^{11} - \ddot{M}^{22} \right) \\ &= -\frac{4}{rc^4} (G\mathcal{M})^{5/3} \omega^{2/3} \cos(2\omega t) , \\ h_\times(t) &= \frac{2G}{rc^4} \ddot{M}^{12} \end{aligned}$$

$$= -\frac{4}{rc^4} (GM)^{5/3} \omega^{2/3} \sin(2\omega t).$$

We see here that the difference in the amplitude of the waveforms for the different polarizations is a simple phase. However, this is only the case because we made the simplifying assumption that the angular momentum of the system is parallel to the observed propagation direction \hat{z} . To be more general we could also consider the case where the angular momentum of the system is offset from \hat{z} by an inclination ι , that is, where

$$\begin{aligned} \mathbf{r}_1 &= \frac{\mu}{m_1} R (\cos(\omega t), \cos(\iota) \sin(\omega t), \sin(\iota) \sin(\omega t)) , \\ \mathbf{r}_2 &= -\frac{\mu}{m_2} R (\cos(\omega t), \cos(\iota) \sin(\omega t), \sin(\iota) \sin(\omega t)) , \end{aligned} \quad (2.18)$$

and by repeating the above process we get

$$\begin{aligned} h_+(t) &= -\frac{4}{rc^4} (GM)^{5/3} \omega^{2/3} \left[\frac{1 + \cos^2(\iota)}{2} \right] \cos(2\omega t) , \\ h_\times(t) &= -\frac{4}{rc^4} (GM)^{5/3} \omega^{2/3} [\cos(\iota)] \sin(2\omega t) . \end{aligned} \quad (2.19)$$

Including the inclination in this way shows that, as long as we are sensitive to both polarizations separately, we can infer the inclination of a binary system by analyzing the relative amplitudes of h_+ and h_\times . Nonetheless, we can see that, to leading order, the radiation emitted by a binary system in this formalism is monochromatic and has a set frequency.

However, by radiating gravitational waves, the system must necessarily be losing energy. The power radiated by systems as the one treated above is given by:

$$\begin{aligned} P &= \frac{dE}{dt} = \frac{r^2 c^3}{32\pi G} \int d\Omega \sum_{ij} \langle (\partial_t h_{ij}^{TT})^2 \rangle \\ &= \frac{r^2 c^3}{16\pi G} \int d\Omega \left[\langle (\partial_t h_\times^{TT})^2 \rangle + \langle (\partial_t h_+^{TT})^2 \rangle \right] , \end{aligned}$$

where the angular brackets represent the time averaging of the inner term over time (so that P is the *averaged* energy loss over some arbitrary number of cycles). The angular distribution is then

$$\frac{dP}{d\Omega} = \frac{r^2 c^3}{16\pi G} \cdot \left(\frac{4}{rc^4} (GM)^{5/3} \omega^{2/3} \right)^2 \cdot g(\iota) \left[\langle (\partial_t \cos(2\omega t))^2 \rangle + \langle (\partial_t \sin(2\omega t))^2 \rangle \right]$$

$$\begin{aligned}
&= \frac{1}{c^5 \pi G} \cdot (GM\omega)^{10/3} \cdot g(\iota) \cdot [2 + 2] \\
&= \frac{4}{\pi G c^5} \cdot (GM\omega)^{10/3} \cdot g(\iota),
\end{aligned}$$

with $g(\iota) = [1/2 + \cos^2(\iota)/2]^2 + \cos^2(\iota)$. Integrating $g(\iota)$ over the solid angle we have $\int d\Omega g(\iota) = 2\pi \int d\cos(\iota) g(\iota) = 8\pi/5$, and so the total radiation power is

$$P = \frac{32}{5} \cdot \frac{(GM\omega)^{10/3}}{Gc^5}.$$

Given that the energy of the orbit for a binary system is $E_{orb} = -Gm_1m_2/2R$, and $\frac{dE_{orb}}{dt} = Gm_1m_2\dot{R}/2R^2$. Then, to compensate for GW radiation,

$$\frac{Gm_1m_2}{2R^2}\dot{R} = -\frac{32}{5} \cdot \frac{(GM\omega)^{10/3}}{Gc^5}, \quad (2.20)$$

which can only be accomplished by $\dot{R} < 0$. Using (2.16), then

$$\begin{aligned}
\dot{R} &= -\frac{2}{3}R\frac{\dot{\omega}}{\omega} \\
&= -\frac{2}{3}(Gm\omega)^{1/3}\frac{\dot{\omega}}{\omega^2}
\end{aligned} \quad (2.21)$$

implies $\dot{\omega} > 0$. The emission of gravitational waves means the component masses of the binary system grow closer and closer together over time, while the frequency of orbital motion increases. We can recast (2.20) as

$$-\frac{Gm_1m_2}{3(G\frac{m}{\omega^2})^{2/3}}(Gm\omega)^{1/3}\frac{\dot{\omega}}{\omega^2} = -\frac{32}{5} \cdot \frac{(GM\omega)^{10/3}}{Gc^5},$$

and we can solve for $\dot{\omega}$ to obtain

$$\dot{\omega} = \frac{96}{5c^5}(GM)^{5/3}\omega^{11/3}, \quad (2.22)$$

which has solutions of the type

$$\omega(t) = \left(\frac{5}{256} \frac{c^5}{(t_{coal} - t)(GM)^{5/3}} \right)^{3/8} \quad (2.23)$$

where t_{coal} is some constant we interpret as the time at which the binary system coalesces.




Figure 2.1: Evolution of the angular frequency $\omega(t)$ over time for a binary system with chirp mass $\mathcal{M} = 20 M_{\odot}$.

We see that at $t = t_{coal}$, $\omega(t)$ actually diverges, as plotted in Figure 2.1. However, from equation (2.16), this would mean $R(t = t_{coal}) = 0$. Since we do not expect the bodies in our problem to actually be point masses in a real physical system, we can assume that the two bodies collide and merge after reaching some critical separation $R_{crit} > 0$, after which (2.23) is no longer valid. Note that from (2.21) we have $\dot{R}/R = -(2\dot{\omega}/3\omega)$, which we can integrate in time to obtain

$$R(t) = R_0 \left(\frac{t_{coal} - t}{t_{coal} - t_0} \right)^{1/4}$$

where R_0 is the initial value of R at some initial time t_0 . We can now adjust our coordinates in (2.18) to account for this time dependence:

$$\begin{aligned} \mathbf{r}_1 &= \frac{\mu}{m_1} R(t) (\cos(\Phi(t)), \cos(\iota) \sin(\Phi(t)), \sin(\iota) \sin(\Phi(t))) , \\ \mathbf{r}_2 &= -\frac{\mu}{m_2} R(t) (\cos(\Phi(t)), \cos(\iota) \sin(\Phi(t)), \sin(\iota) \sin(\Phi(t))) , \end{aligned}$$

with

$$\Phi(t) = \int_{t_0}^t \mathbf{d}\tau \omega(\tau) = -2 \left[c^3 \frac{t - t_{coal}}{5G\mathcal{M}} \right]^{5/8} + \Phi_0 ,$$

where Φ_0 is an integration constant.

We can in principle now calculate the time derivatives of the quadrupole mass moment for the time dependent case in order to build the waveforms h_+ and h_\times . However, we will first argue that the contributions from the derivatives of $R(t)$ and $\omega(t)$ are negligible and can thus be ignored to leading order. This is because, under the test mass limit, where one of the masses is much lighter than the other, the radial equations of motion under the geometry of the Schwarzschild solution tell us that in the equations of motion for the radial coordinate there is an effective gravitational potential of the form

$$U_{eff} = \left(1 - \frac{R_s}{R}\right) \left(c^2 + \frac{\Lambda}{R}\right),$$

where $R_s = 2Gm/c^2$, with m being the total mass of the system, is approximately the Schwarzschild radius of the heavier object, and Λ is the specific angular momentum of the system. This potential has a minimum at

$$R = \frac{\Lambda^2 + \sqrt{\Lambda^4 - 3c^2 R_s^2 \Lambda^2}}{c^2 R_s}.$$

This value is minimized when the expression under the square root is null, which means setting $\Lambda = \sqrt{3}cR_s$. Substituting back in the expression for R , we obtain the minimum separation beyond which no stable circular orbits are allowed, that is, the Innermost Stable Circular Orbit (ISCO), expressed by

$$R_{ISCO} = 3R_s = \frac{6Gm}{c^2}.$$

This in turn implies a maximum frequency,

$$\omega_{ISCO} = \frac{c^3}{2\sqrt{54}Gm}.$$

If we abandon the test mass limit and consider a binary system of comparable masses, the above equations will require some further treatment [35], but this approximation will serve for our purposes. From (2.21) we have that \dot{R} is negligible so long as $\dot{\omega} \ll \omega^2$ which from (2.22) is equivalent to the condition $\omega \ll c^3/2GM$. Since this value is larger than ω_{ISCO} , after which the inspiral regime we have been working on is no longer valid, we can neglect the effects of \dot{R} , and similarly of $\dot{\omega}$. Because of this, we simply have to substitute the ωt factor in the argument of the trigonometric functions in (2.19) with $\Phi(t)$, as well as substituting ω




Figure 2.2: Plot of $h(t)$ for gravitational waves emitted by a system with $\mathcal{M} = 20 M_{\odot}$, $r = 0.1$ Gpc and $\iota = \pi/2$.

with our expression for $\omega(t)$, and we obtain

$$\begin{aligned}
 h_{+}(t) &= -\frac{GM}{rc^3} \left(5 \frac{cGM}{(t_{\text{coal}} - t)} \right)^{1/4} \left[\frac{1 + \cos^2(\iota)}{2} \right] \cos(2\Phi(t)), \\
 h_{\times}(t) &= -\frac{GM}{rc^3} \left(5 \frac{cGM}{(t_{\text{coal}} - t)} \right)^{1/4} [\cos(\iota)] \sin(2\Phi(t)).
 \end{aligned}
 \tag{2.24}$$

These waveforms increase in amplitude and frequency with time, leading to the characteristic “chirp” signal of gravitational waves as seen in Figure 2.2. This type of signature is what we will be exploring in the experimental part of this dissertation.

2.4 Post Newtonian expansion and higher order contributions

The formalism we just developed is a mere approximation. In particular, a strong assumption that was made was that the background spacetime to our perturbations was flat, and independent of the motions of the GW source. On the whole, this is not a reasonable assumption to make and is only applicable to systems whose gravitational fields near the source are small,

which is a particularly bold assumption for binary systems. A more common method used is the Post-Newtonian (PN) formalism, which expands on powers of $x = (v/c)^2$, is applicable to a wider range of systems. In these systems, at higher orders, different physical factors come into effect. Of particular interest to us will be the effect of the angular momentum of the component bodies of a binary, or spin. This spin is a vectorial quantity, and it can be projected into the orbital angular momentum in order to obtain parallel and perpendicular components, which have distinct effects on the GW waveforms. The parallel component, if aligned (anti-aligned) with the orbital spin, will cause the binary system to merge at smaller (larger) separation distances. This is an effect of order 1.5PN, that is, it contains a factor $(v/c)^3$, and will be of particular interest to us when dealing with experimental data. The perpendicular component of the spin, which is not going to be explored further in this work, will cause the orbital angular momentum to precess with respect to the total angular momentum, which will cause modulations on the amplitude of $h(t)$. A thorough review of the PN formalism, covering the basics as well as spinning and eccentric cases, can be found in [36].

3 Experimental setup

3.1 Michelson Interferometers

Modern gravitational wave detector designs are based on Michelson-type interferometers, instruments which are extraordinarily accurate at measuring changes in the travel time of light in its arms. In this chapter we take a quick dive into the inner workings of this type of devices. We start by pointing to the diagrammatic representation of a basic Michelson interferometer in Figure 3.1: a monochromatic light source¹ fires a beam of light towards a beam splitter, which splits the beams into two with equal probability. The two resulting beams travel through the two optical paths at the orthogonal arms of the interferometer, which should be of equal length, then reflected back into the beam splitter. Here each of the two beams splits again, combining into two different paths: in one path, part of the reflected light of the horizontal arm passes through the splitter while part of the reflected light from the vertical arm is reflected at the splitter, and the two beams combine following a path to the detector. A second combined beam goes towards the source and is irrelevant for practical purposes, and often (as in our case) is not even diagrammatically shown. Depending on the difference between the lengths of the optical paths, the light on the detector path will suffer interference, modulating the amplitude of the power registered by the detector. We will now show the mathematical description of this process.

Consider our monochromatic source to be a laser with frequency ω_L . Equivalently, we have a laser with wavelength $\lambda_L = 2\pi/\omega_L$ and wavenumber $k_L = \omega_L/c$. The electric field generated is then:

$$E(t, \mathbf{x}) = E_0 e^{(-i\omega_L t + i\mathbf{k}_L \cdot \mathbf{x})}.$$

Take the origin of our coordinate system to be at the beam splitter, so that here $\mathbf{k}_l \cdot \mathbf{x} = 0$. The overall phase of the field is conserved during free propagation. Transmission through the beam splitter introduces a factor of $1/\sqrt{2}$ in the amplitude, while reflection contributes with another factor of $1/\sqrt{2}$, as well as a phase π (that is, an overall (-1) factor) due to reflection on one of the sides.² The reflections at the mirrors at the end of the arms also contribute a phase of π . As such, if we consider the final contribution of the beam traveling through the horizontal arm of length L_x as E_x and the contribution of the beam traveling

¹While monochromatic sources have always been the ideal source in theory, it was not uncommon for early researchers, Michelson and Morley included, to use (quasi-)monochromatic light to set up the equipment and use white light for the measurements themselves [37].

²The phase factors happens only on one of the sides depending on the diffraction coefficient of the dielectric coating on the beam splitter.

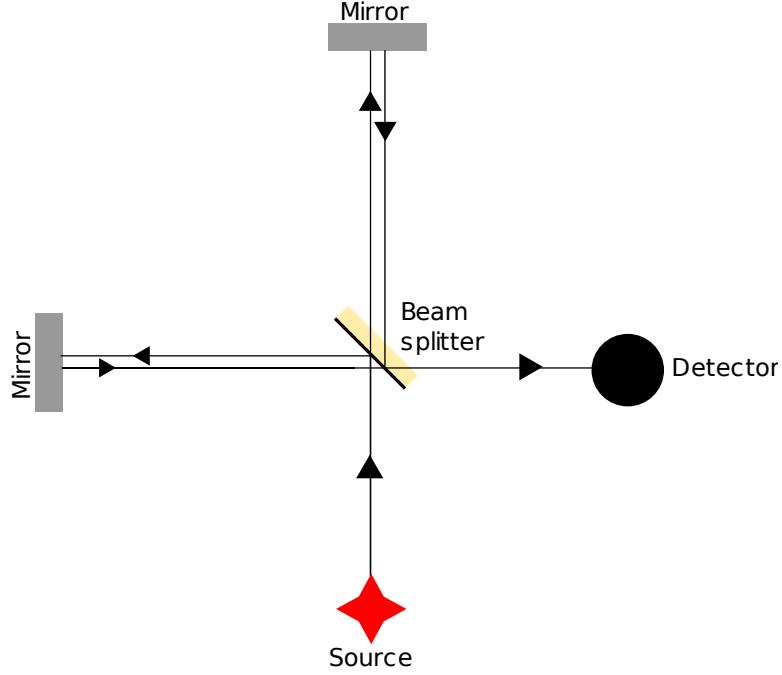


Figure 3.1: The layout for a basic Michelson interferometer. The yellow part of the beam splitter represents a dielectric coating.

through the vertical arm of length L_y as E_y , and defining the initial times $t_0^x = t - 2L_x/c$ and $t_0^y = t - 2L_y/c$, we have

$$E^x = -\frac{1}{2}E_0e^{(-i\omega_L t + ik_L L_x)},$$

$$E^y = \frac{1}{2}E_0e^{(-i\omega_L t + ik_L L_y)},$$

and so the electric field at the time of recombination is:

$$\begin{aligned} E^r &= E^x + E^y \\ &= \frac{1}{2}E_0e^{-it\omega_L} (e^{ik_L L_y} - e^{iL_x k_L}) \\ &= iE_0e^{i[k_L(L_y+L_x)-t\omega_L]} \sin(k_L(L_y - L_x)). \end{aligned}$$

The intensity measured at the detector is then proportional to:

$$|E^r|^2 = E_0^2 \sin^2(k_L(L_y - L_x)),$$

meaning that any deviation in the length of the optical paths of the arms will be translated

to a change in the intensity registered at the detector. Since gravitational waves deform spacetime and change the geodesics along their paths, the passing of a GW through an interferometer would change the travel time of light along the arms, effectively stretching and shrinking them, which would leave a detectable signature.

3.1.1 Interaction of a gravitational wave with a Michelson interferometer

We saw in 2.2.1 how gravitational waves in the TT gauge would interact with a couple of test masses. We will now develop that further and see how these ripples in spacetime would affect the electric field arriving at the detector. Let us consider again the position of the beam splitter as the origin of our cartesian coordinate system, and define the position of the mirrors

$$\begin{aligned}\mathbf{r}_x &= (L_x, 0, 0) , \\ \mathbf{r}_y &= (0, L_y, 0) .\end{aligned}$$

It is important here to raise the point that, as we saw before, the effect of the gravitational waves in the TT gauge does *not* manifest itself in the coordinates, but rather it will affect the spacetime interval and, as such, the propagation of light between two points.

We will start by considering the simpler case of a wave which has only the plus polarization (such as one emitted from a binary system with inclination $\iota = \pi/2$), and as usual we will take it to be propagating in the z direction. We will neglect for now the effect of orbital decay, and we will for convenience use the GW frequency $\omega_{gw} = 2\omega$ rather than the orbital frequency ω . Then, in the plane of the interferometer arms, from 2.19 we get

$$h_+(t) = h_0 \cos(\omega_{gw}t) .$$

To obtain the relevant spacetime interval we can look at equation (2.10) to obtain

$$ds^2 = -c^2 dt^2 + [1 + h_+(t)] dx^2 + [1 - h_+(t)] dy^2 + dz^2 . \quad (3.1)$$

If we consider the x arm, and knowing that photons follow null geodesics (so $ds^2 = 0$), we have

$$dx^2 = \frac{c^2 dt^2}{[1 + h_+(t)]} .$$

We can reasonably approximate $|h_0| \ll 1$ to obtain

$$\begin{aligned} dx &= \pm c dt \sqrt{\frac{1}{[1 + h_+(t)]}} \\ &\approx \pm c dt \left[1 - \frac{1}{2} h_+(t) \right]. \end{aligned} \quad (3.2)$$

We have two solutions, the positive one corresponding to the travel from the beam splitter to the mirror, and the negative corresponding to the return trip. Let us now consider a photon leaving the beam splitter at time t_0 . If that is the case, it will reach the mirror at \mathbf{r}_x at some time t_1 which can be obtained by integrating the positive solution of equation (3.2)

$$\begin{aligned} \int_0^{L_x} dx &= c \int_{t_0}^{t_1} dt \left[1 - \frac{1}{2} h_+(t) \right] \\ \rightarrow L_x &= c(t_1 - t_0) - \frac{c}{2} \int_{t_0}^{t_1} dt h_+(t), \end{aligned}$$

and at the same time, in the return trip from the mirror to the beam splitter, arriving at some time t_2 , we have

$$\begin{aligned} \int_{L_x}^0 dx &= -c \int_{t_1}^{t_2} dt \left[1 - \frac{1}{2} h_+(t) \right] \\ \rightarrow L_x &= c(t_2 - t_1) - \frac{c}{2} \int_{t_1}^{t_2} dt h_+(t), \end{aligned}$$

so we can sum both paths and get

$$2L_x = c(t_1 - t_0) + c(t_2 - t_1) - \frac{c}{2} \int_{t_0}^{t_1} dt h_+(t) - \frac{c}{2} \int_{t_1}^{t_2} dt h_+(t),$$

which with some minor rearrangement gives us

$$(t_2 - t_0) = \frac{2L_x}{c} + \frac{1}{2} \int_{t_0}^{t_2} dt h_+(t). \quad (3.3)$$

We see then that in the absence of gravitational waves the round trip time of light in the TT gauge is simply $2L_x/c$, that is, the total distance traveled over the speed of said light. However, the presence of gravitational waves introduces a small correction given by the integral in equation (3.3). Since from equation (3.2) we are already neglecting any contributions of

order $O(h_0^2)$ and above, we can approximate $t_2 \approx t_0 + 2L_x/c$ in the upper limit of the integral. This limits our calculation to first order, and we get

$$\begin{aligned}
t_2 - t_0 &= \frac{2L_x}{c} + \frac{1}{2} \int_{t_0}^{t_0 + \frac{2L_x}{c}} dt h_0 \cos(\omega_{gw} t) \\
&= \frac{2L_x}{c} + \frac{h_0}{2\omega_{gw}} [\sin(\omega_{gw} t)]_{t_0}^{t_0 + \frac{2L_x}{c}} \\
&= \frac{2L_x}{c} + \frac{h_0}{2\omega_{gw}} \left[\sin\left(\omega_{gw} \left[t_0 + \frac{2L_x}{c}\right]\right) - \sin\left(\omega_{gw} \left[t_0 + \frac{2L_x}{c}\right]\right) \right], \quad (3.4)
\end{aligned}$$

and using the identity $\sin(\alpha + 2\beta) - \sin(\alpha) = 2 \sin(\beta) \cos(\alpha + \beta)$,

$$\begin{aligned}
(t_2 - t_0) &= \frac{2L_x}{c} + \frac{1}{\omega_{gw}} h_0 \cos\left(\omega_{gw} \left[t_0 + \frac{L_x}{c}\right]\right) \sin\left(\omega_{gw} \frac{L_x}{c}\right) \\
&= \frac{2L_x}{c} + \frac{1}{\omega_{gw}} h \left(t_0 + \frac{L_x}{c}\right) \sin\left(\omega_{gw} \frac{L_x}{c}\right).
\end{aligned}$$

We can multiply and divide the correction term by L_x/c to rewrite the above equation as

$$(t_2 - t_0) = \frac{2L_x}{c} + \frac{L_x}{c} h \left(t_0 + \frac{L_x}{c}\right) \text{sinc}\left(\omega_{gw} \frac{L_x}{c}\right), \quad (3.5)$$

where $\text{sinc}(x) = \sin(x)/x$, which goes to one as $x \rightarrow 0$. We can see then, that as long as $\omega_{gw} L_x \ll c$, we can again approximate the result as follows

$$(t_2 - t_0) = \frac{2L_x}{c} + \frac{L_x}{c} h \left(t_0 + \frac{L_x}{c}\right). \quad (3.6)$$

On the other hand, if we were to consider the extremely high frequency or extremely large arm cases then as $\lim_{x \rightarrow \infty} \text{sinc}(x) = 0$ we conclude that the correction term disappears. Intuitively, we can think of this as all the contributions from the peaks (positive) and troughs (negative) of the gravitational wave canceling out over space.

The calculations for the vertical y arm are similar, only as we can see in equation (3.1) the sign of $h(t)$ is inverted. We then get

$$(t_2 - t_0) = \frac{2L_y}{c} - \frac{L_y}{c} h \left(t_0 + \frac{L_y}{c}\right) \text{sinc}\left(\omega_{gw} \frac{L_y}{c}\right). \quad (3.7)$$

In practice we want to calculate what we would measure at the *output* of the beam splitter at some observation time t . We should then consider, for generality, that the contributions for

both arms may have different initial times. Inverting equations (3.6) and (3.7) to first order in h_0 , and taking into account that for arm i we have $t_0^i + L_i/c \approx t - 2L_i/c + L_i/c = t - L_i/c$, we obtain the initial times for each arm:

$$t_0^x = t - \frac{2L_x}{c} - \frac{L_x}{c} h \left(t - \frac{L_x}{c} \right) \text{sinc} \left(\omega_{gw} \frac{L_x}{c} \right), \quad (3.8)$$

$$t_0^y = t - \frac{2L_y}{c} + \frac{L_y}{c} h \left(t - \frac{L_y}{c} \right) \text{sinc} \left(\omega_{gw} \frac{L_y}{c} \right). \quad (3.9)$$

Now, as we saw before, the phase of the electric field at time t ought to be the same as the initial phase, save for a factor of $\pm 1/2$, so:

$$\begin{aligned} E^x(t) &= -\frac{1}{2} E_0 e^{-i\omega_L t_0^x} \\ &= -\frac{1}{2} E_0 e^{-i\omega_L \left(t - \frac{2L_x}{c} \right) + i\Delta\phi^x(t)}, \end{aligned} \quad (3.10)$$

with

$$\Delta\phi^x(t) = \frac{\omega_L L_x}{c} h \left(t - \frac{L_x}{c} \right) \text{sinc} \left(\omega_{gw} \frac{L_x}{c} \right). \quad (3.11)$$

Similarly, for the y beam we have

$$\begin{aligned} E^y(t) &= \frac{1}{2} E_0 e^{-i\omega_L t_0^y} \\ &= \frac{1}{2} E_0 e^{-i\omega_L \left(t - \frac{2L_y}{c} \right) + i\Delta\phi^y(t)}, \end{aligned} \quad (3.12)$$

where

$$\Delta\phi^y(t) = -\frac{\omega_L L_y}{c} h \left(t - \frac{L_y}{c} \right) \text{sinc} \left(\omega_{gw} \frac{L_y}{c} \right). \quad (3.13)$$

In general we expect $L_x \approx L_y$, in order to cancel noises common to both arms. With that, we can substitute L_x and L_y in equations (3.11) and (3.13) with $L = (L_x + L_y)/2$, so that $\Delta\phi^y(t) = -\Delta\phi^x(t)$, though we will still consider L_x and L_y as distinct outside of these phase shifts. We can, however, rewrite

$$\begin{aligned} 2L_x &= 2L + (L_x - L_y), \\ 2L_y &= 2L - (L_x - L_y). \end{aligned}$$

and therefore

$$E^x(t) = -\frac{1}{2}E_0 e^{-i\omega_L(t-\frac{2L}{c})+i\phi_0+i\Delta\phi^x(t)},$$

$$E^y(t) = \frac{1}{2}E_0 e^{-i\omega_L(t-\frac{2L}{c})+i\phi_0-i\Delta\phi^x(t)},$$

where $\phi_0 = k_L(L_x - L_y)$. The total recombined electric field at the output of the beam splitter is then

$$E^r(t) = E^x(t) + E^y(t)$$

$$= -iE_0 e^{-i\omega_L(t-\frac{2L}{c})} \sin[\phi_0 + \Delta\phi^x(t)].$$

In the $\omega_{gw}L/c$ limit we have

$$\phi_0 + \Delta\phi^x(t) = k_L \left[(L_x - L_y) + Lh \left(t - \frac{L_x}{c} \right) \right],$$

which contextualizes the phase shift resulting from the effect of the gravitational waves as equivalent to a change in $(L_x - L_y)$. With all this, we have finally the power of the electric field at the detector as

$$P = P_0 \sin^2[\phi_0 + \Delta\phi^x(t)]$$

$$= \frac{P_0}{2} (1 - \cos[2\phi_0 + 2\Delta\phi^x(t)]), \quad (3.14)$$

which shows us that if we want the effect of the gravitational waves to be as noticeable as possible then we will want to do what we can to maximize $\Delta\phi^x(t)$.³ Looking back at equation (3.11) we see that the dependence on L is $(\omega_L L/c) \text{sinc}(\omega_{gw}L/c) = (\omega_L/\omega_{gw}) \sin(\omega_{gw}L/c)$. To maximize the sin factor we then want

$$\omega_{gw} \frac{L}{c} = \frac{\pi}{2},$$

or equivalently

$$L = \frac{c}{4f_{gw}} \simeq \frac{75000 \text{ km}}{f_{gw}}.$$

This means that if we wanted to be sensitive to waves of frequency, say, 100 Hz, we would need an arm around 750 km long. Arms with such dimensions are not practically or economically

³Note that for equal length arms, in the absence of any physical effect we have $P = 0$. This is relevant because any disturbance in the laser's power source (which could be modeled by giving P_0 a time dependence) will not be picked up.




Figure 3.2: Layout of a Michelson interferometer with Fabry-Perot cavities. Note how the length L now refers to the length of the cavity, and we take the distance traveled outside the cavity to be negligible.

reasonable to build. A way around this is to use Fabry-Perot (FP) cavities, using coupled mirrors on each arm to make the laser light bounce back and forth and thus increase its travel time. A diagrammatic representation is given in Figure 3.2. For full details on the physical processes behind the functioning of FP cavities we direct the reader to section 9.2 of [38]. The bottom line, however, is that the application of Fabry-Perot cavities in the arms of a Michelson-type GW interferometer will cause photons to, on average, be “stored” in the arms for some storage time

$$\tau_s \simeq \frac{L\mathcal{F}}{c\pi},$$

where L is now the length of the FP cavities (we assume the whole length of what we previously considered arm is converted to a cavity) and \mathcal{F} is a numerical factor known as the *finesse*, given by

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2},$$

where, for a given arm, r_1 is the reflectivity of the first mirror in the cavity and r_2 is the reflectivity of the second mirror. We can also define the pole frequency

$$f_p = \frac{1}{4\pi\tau_s} \simeq \frac{c}{4\mathcal{F}L}$$

As long as $\omega_{gw}L_x/c \ll 1$ and $\mathcal{F} \sim (\omega_{gw}L_x/c)^{-1}$, this will ultimately contribute to the phase shift seen in equation (3.11) as follows:

$$\Delta\phi^x(t) \simeq h_0 \frac{4\mathcal{F}}{\pi} k_L L \frac{1}{\sqrt{1 + \left(\frac{f_{gw}}{f_p}\right)^2}} \cos \left[\omega_{gw} \left(t - \frac{L}{c} \right) \right].$$

Note that for $f_{gw} \gg f_p$ we have

$$\Delta\phi^x(t) \simeq h_0 \frac{4\mathcal{F}}{\pi} k_L L \frac{f_p}{f_{gw}} = \frac{h_0 k_L c}{\pi f_{gw}},$$

and we see that the sensitivity of the detector decreases linearly with f_{gw} . This degradation is compounded by the $\text{sinc}(\omega_{gw}L/c)$ factor, which we can no longer approximate as unity when f_{gw} increases and so deteriorates the response further. This highlights how carefully these instruments must be tuned for proper functioning.

3.1.2 Geometrical dependence of a GW interferometer detection

Another relevant effect to cover is the effect of the source's sky position on the detectors' sensitivity. Let us consider a reference frame centered at the beam splitter and with arms along the x and y axes, and a second reference frame in the sky as displayed in 3.3. If we take the polarization matrices in h_+ and h_\times to be referent to the sky coordinate system, then the strain of the interferometer is given by

$$\frac{\delta L(t)}{L} = F_+(\theta, \phi, \psi) h_+(t) + F_\times(\theta, \phi, \psi) h_\times(t), \quad (3.15)$$

where F_+ and F_\times are known as the *antenna pattern response functions* for each polarization, which can be seen as projecting the effect of the gravitational waves onto the detector frame [39]. Using the geometry of Figure 3.3 it can be shown that

$$F_+ = \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \cos 2\psi - \cos \theta \sin 2\phi \sin 2\psi, \quad (3.16)$$

$$F_\times = \frac{1}{2} (1 + \cos^2 \theta) \cos 2\phi \sin 2\psi + \cos \theta \sin 2\phi \cos 2\psi. \quad (3.17)$$

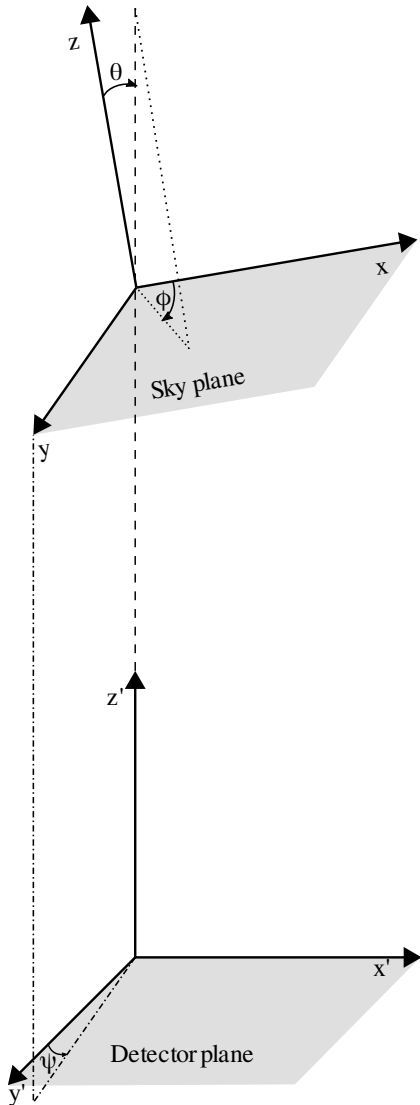


Figure 3.3: Example of the geometry for GW detection.

Note that the maximum value for these functions is 1. For a certain sky position, we can then calculate the *antenna power pattern*

$$P(\theta, \phi) = F_+(\theta, \phi, \psi)^2 + F_\times(\theta, \phi, \psi)^2, \quad (3.18)$$

which is independent from the angle ψ . This gives us the sensitivity of the detector as a whole to the sky position of the source. However, we would like to extend this to multiple-detector networks, as in the real case of the LIGO-Hanford, LIGO-Livingston and Virgo detectors, which have been working in tandem since 2016. The first thing to take care of in order to achieve this generalization is to use a common coordinate system for all detectors, such as earth-based spherical coordinates. An in depth treatment of this problem, taking into account contributions from such factors as noise correlations in detector networks and the rotation of the earth can be found in references [40, 41]. We will limit ourselves to the simpler (yet reasonable) case of uncorrelated detector noise and constant time, under which the normalized *network antenna power pattern* is given simply by

$$P_N(\theta, \phi) = \sum_k (F_{+,k}^2 + F_{\times,k}^2), \quad (3.19)$$

where the subscript k stands for the detector considered for each term. With this quantity we can check the full network's sensitivity to a certain sky position, at a certain time during earth's rotation (we expect the network's coverage map to be cyclical over a full rotation). An example of this pattern is given in Figure 3.4. Note that we seem to find a reflection symmetry in the sky position. This is because, with 3 detectors, a naive triangulation of a signal originating from sky position \mathbf{r} is indistinguishable from a signal originating at $-\mathbf{r}$, though requirement amplitude and phase consistency can break this degeneracy in some cases [42].

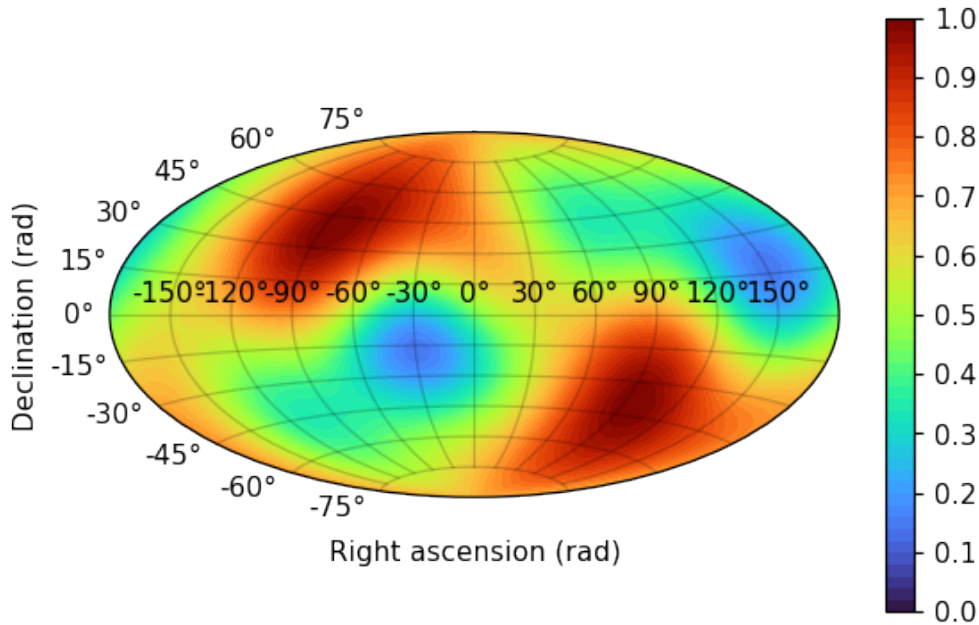


Figure 3.4: Map of the network antenna power pattern for the advanced LIGO/Virgo network at gps time $(t \bmod 24 \text{ h}) = 9742 \text{ s}$.

3.2 Noise in the advanced LIGO and advanced Virgo detectors

The previous subsection explores an idealized version of gravitational wave interferometry. We turn now to explore some considerations arising from the use of real-world detectors in gravitational wave astronomy. Current GW interferometers are essentially Michelson-type interferometers as discussed before, but what makes these instruments marvels of modern engineering is the tremendous sensitivity they can achieve despite all kinds of environmental difficulties. In this subsection we will qualitatively cover some of the more significant noise sources that challenge the detection of GWs (see Figure 3.5) and the systems in place that have allowed the LVC detectors to overcome them.

3.2.1 Seismic noise

Seismic noise comes from ground motion due to the internal dynamics of the earth (or environmental disturbances such as winds, ocean waves and human activity) being transmitted to the components of the interferometer, with the largest influence on the signal being the effect on the device's test masses and beam splitter. The Advanced LIGO and Advanced Virgo observatories are equipped with highly advanced hybrid mechanical isolation systems (passively isolating as many perturbations as possible, while using sensors and actuators

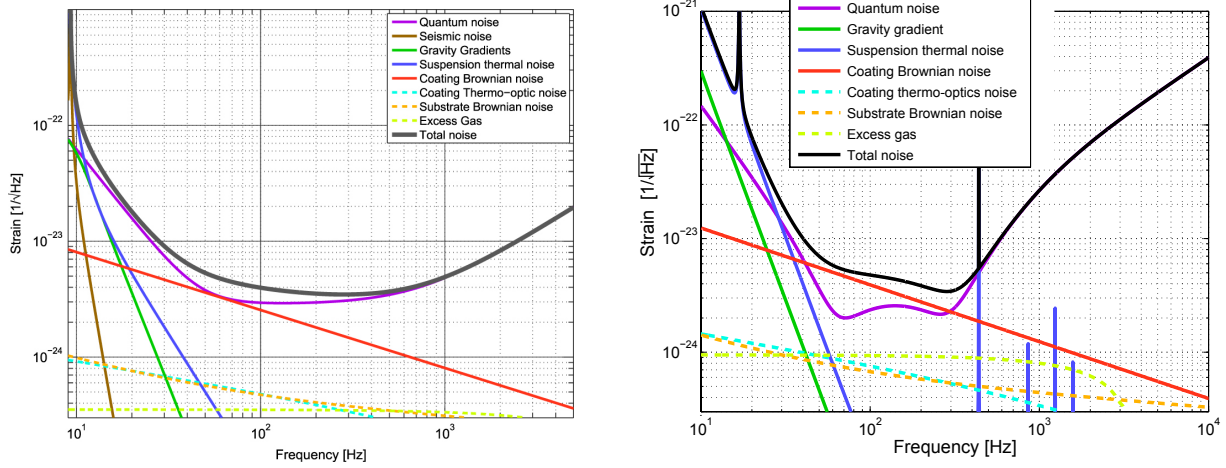


Figure 3.5: **Left:** Advanced LIGO design sensitivity curves from [30]. **Right:** Advanced Virgo design sensitivity curves from [43].

to actively counter inevitable motions) in order to combat this, which manage to significantly mitigate seismic noise, especially for frequencies higher than 20 Hz [30, 31, 44, 45]. A representation of these can be found in Figure 3.6.

3.2.2 Gravity gradients

Gravity gradient noise comes from density perturbations in the earth caused by seismic activity. These subtle fluctuations produce changes in the local gravitational forces that the test masses are subjected to. This type of noise is fundamental and is a limiting factor for both LIGO and Virgo's strain sensitivity, with proposals for mitigation dependent on increasingly accurate modelling of these perturbations [46, 47]. However, the current generation of detectors is still far from this limit [48].

3.2.3 Quantum noise

Quantum noise refers to artifacts in the interferometer output related to the quantum nature of the electromagnetic fields involved in measurements. There are a number of phenomena that contribute to this type of noise: on one hand, there is the issue of the intrinsic quantum mechanical uncertainty of the mirrors' position and momenta and the resulting back-action phenomena, which impose a limit on the minimum measurement error. On the other hand, radiation pressure noise is induced by quantum vacuum fluctuations at the unused input port of the beam splitter. Finally, uncertainties regarding the number of photons at the interferometer output originate a phenomenon known as shot noise. Techniques to mitigate

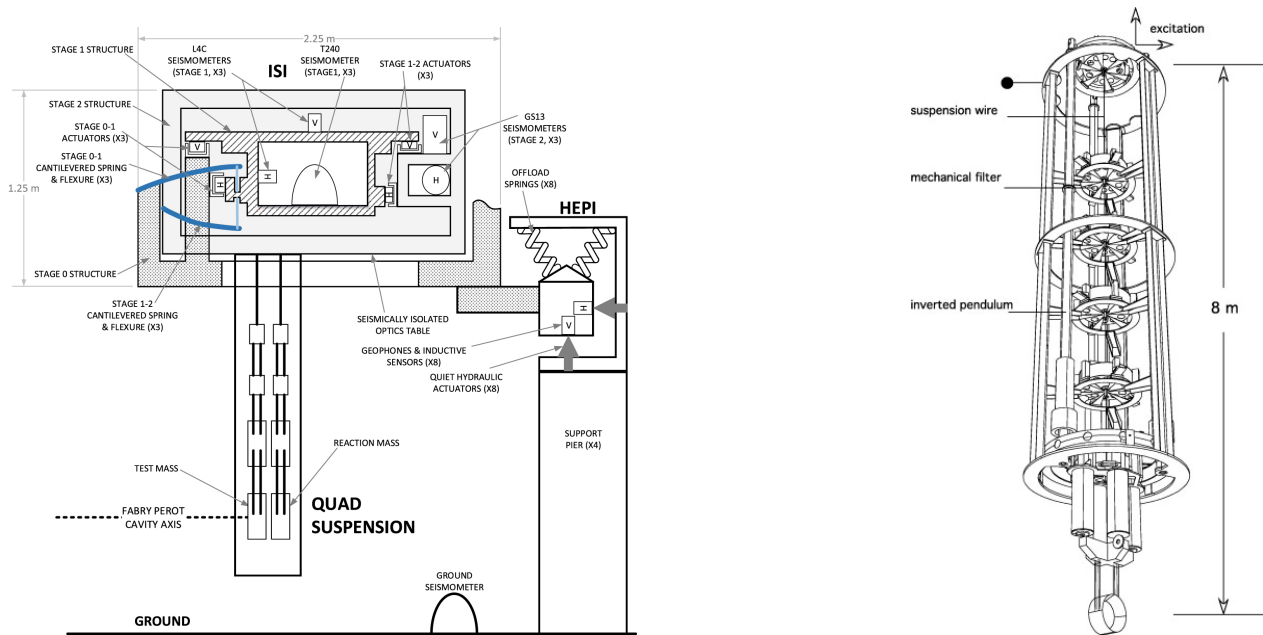


Figure 3.6: Seismic isolation systems in gravitational wave interferometers. **Left:** the seismic isolation system used in LIGO, from [30] **Right:** Virgo’s Superattenuator, from [45]

these perturbations usually involve the usage of squeezed states of light. [49, 50, 51]

3.2.4 Suspension thermal noise

Suspension thermal noise is, as the name indicates, the noise resulting from random motion of the mirrors’ suspension due to thermal effects. The current generation of GW detectors has significantly minimized this type of noise relative to previous iterations by using fused silica suspensions in a monolithic arrangement (rather than thin steel wires in a “cradle-like” arrangement), obtaining increased sensitivity by up to a factor of 10 for the target range of frequencies [52, 53].

3.2.5 Coating brownian noise

Despite having exceptional optical properties, the silica-tantala ($\text{SiO}_2\text{-Ta}_2\text{O}_5$) multilayer coatings on LIGO and Virgo mirrors give rise to significant thermal (brownian) noise dominantly due to mechanical losses in the tantala layers. The mitigation of this effect is a subject of on-going research and may be achieved through a combination of doping optimization, annealing techniques and alternative coating materials [30, 54, 55].

3.2.6 Glitches

Glitches are transient (i.e. short-lasting) noise signals of instrumental or environmental origin that can interfere with the regular functioning of the detector. Besides the obvious way in which additional noise can decrease sensitivity to astrophysical GW signals, some glitch morphologies (such as blip glitches) may mimic the signatures of GW events, sometimes as often as once per minute [56], though this can be mitigated by requiring GW data to be consistent in more than one detector (as non-astrophysical signals should not generate correlated signals at every detector in the network). Recent research has tried to find ways to classify and remove such glitches from strain data using machine learning methods [57, 58, 59]. Furthermore, a citizen science project known as Gravity Spy [60] makes use of human volunteers to characterise and analyse LIGO and Virgo glitches.

4 Machine Learning

Artificial Intelligence (AI) is a scientific term which covers, in a general way, any process of making computers behave in a manner that resembles human cognition. As a field of research, Artificial Intelligence was kickstarted in the mid-1950s, with the conceptualization of the “imitation game” by Alan Turing in 1950 (later to be known as the Turing test) serving as a foreboding overture to the 1956 Dartmouth Summer Research Project on Artificial Intelligence, which springboarded much of the research done in AI for the following two decades. For a long time, the focus of research was on so-called Expert Systems: complex hand-written programs with a large set of rules and facts, often employing brute search to perform a pre-defined, unchangeable task. A mediatic example of such a system is IBM’s Deep Blue, a chess computer that famously and controversially beat then chess world champion Garry Kasparov in 1997 [61]. However, a problem with such systems is that they require the tasks to already have some previously known algorithmic solution so that rules can be implemented, and thus do not perform as well for use cases where variance plays a large part, or where arbitrary abstractions may play a role. For these cases, examples of which include tasks like image classification, or finding unknown patterns in data, a different approach is required.

Machine Learning (ML) is the name given to an implementation of AI which takes strides towards the mimicking, understanding and aiding of human cognitive processes by processing and performing inference on, usually, very large amounts of data. In particular, the term covers algorithms which are capable of automatic self-improvement as they tackle some given problem, having been fed some amount of training data to “learn” from. The field of ML has supercharged applications such as natural language processing (NLP) and image recognition. Deep Learning (DL), an implementation of ML where the data is parsed through several layers of increasing abstraction in a neural network (NN), which is inspired by the structure of the human brain, has become particularly prevalent over the last decade.

ML algorithms can roughly be separated into two categories:

- **Supervised learning**, where, during training, the target outputs (often called “labels”) for input data are fed to the machine, and so the algorithm’s goal is essentially to find some configuration of the machine’s parameters in order to behave as a function mapping the input data to the targets, minimizing some loss function comparing the algorithm output to the target.
- **Unsupervised learning**, where during training there are no target outputs, but rather

the algorithm is asked to find structures and create compact descriptions of the data. Any new data input to the algorithm will be fitted to one of the created groups.

Within these two broader classes there are many subcategories and niche implementations which we do not deem necessary to explore here, as our main focus will be on supervised learning applied to image processing. In this chapter we will explore supervised learning in particular: we shall look at the concept of a neural network, a particular family of NN architectures called Convolutional Neural Networks (CNNs), go over a specific implementation of these networks in the form of Residual Networks (ResNets) applied to classification and regression problems, and finish by exploring a technique to associate a level of uncertainty to a prediction by using Monte-Carlo (MC) dropout. The main references for this chapter are [62, 63].

4.1 Deep learning and neural networks

Neural networks are a type of ML algorithm (or model) architecture that takes inspiration from the structure of the human brain. As the basic unit of the human brain is the neuron, where signals from a number of dendrites are processed in the nucleus and then transmitted through to the axon as the presence (or lack thereof) of an electrochemical action potential, the basic unit of the standard NN is the perceptron, an algorithm that takes a number of inputs, performs a given weighted linear operation on them, and outputs a certain value (usually 1 or 0) depending on whether some condition has been met by the internal operation. This can be expressed as a branched function of the inputs \mathbf{x} and the respective weights \mathbf{w} :

$$g(\mathbf{w} \cdot \mathbf{x}) = \begin{cases} f(\mathbf{w} \cdot \mathbf{x}), & \text{if } \mathbf{w} \cdot \mathbf{x} + b > 0 \\ 0, & \text{otherwise} \end{cases}, \quad (4.1)$$

where $f(\mathbf{w} \cdot \mathbf{x})$ is some chosen activation function.

Neural networks are composed of several interconnected layers of such perceptrons, each of which has its characteristic weights that can be readjusted as the model compares its output to the given targets. The number of layers is typically referred to as the *depth* of the network, while the dimensionality of each layer constitutes the *width*. The connectedness of the layers can vary depending on the architecture chosen: a fully connected neural network (FCNN), for example, has all the neurons in a given layer connected to all the neurons in the preceding and following layer (save for the input and output layers, which are only connected to the following and previous layers, respectively). This is shown in Figure 4.1. The updating of




Figure 4.1: Mock model of an FCNN. Each of the nodes is connected to all the nodes in the previous and following layer, with the exception of the first nodes of the hidden layers, which serve as a constant bias b to the next layer.

deep NN weights is typically done by *backpropagation*. Backpropagation is a way to calculate the gradient (with respect to the weights) of some loss function

$$\mathcal{L}[(w_1, w_2, \dots, w_n), X, Y], \tag{4.2}$$

where w_i is the value of the i -th weight, X is the input data and Y is the target value, by recursively multiplying the gradients of the individual layers by means of the chain rule for differentiation. With the gradient calculated, the position in parameter space of the weights vector can be updated as

$$(w'_1, w'_2, \dots, w'_n) = (w_1, w_2, \dots, w_n) - \alpha \nabla \mathcal{L}, \tag{4.3}$$

where α is a scalar multiplier that defines the *step* size. This is illustrated in Figure 4.2. For simple architectures backpropagation can be seen simply as matrix multiplication, and thus the computational performance of the algorithm when updating weights is dependent only on the hardware's capacity to perform matrix multiplications. If our input data has thousands upon thousands of values, such as we see in the pixels of an image (which is essentially a tensor of size $(\text{width}) \times (\text{height}) \times (\text{channels})$), we may find ourselves not wanting to deal with storing




Figure 4.2: Mock example of a 2-parameter loss function and its gradient. For some set of starting values $(w_1^{(0)}, w_2^{(0)})$, the weights can be updated according to the gradient of the loss function.

and updating quite the same amount of weights, be it for memory reasons or performance reasons. For such cases, an increasingly popular solution is the usage of convolutional neural networks.

4.1.1 Convolutional Neural Networks

CNNs are a type of architecture particularly useful for treating grid-like data, such as images or time-sampled data. In CNNs, instead of assigning weights to each of the input datum, we use a smaller weighted convolutional kernel and apply it to the data some number of times. Consider an input which takes the shape of a (n, m) matrix. We can define our kernel to be a matrix of weights of size (k, l) , then apply the kernel over the image, taking the matrix




Figure 4.3: Mock model of a CNN, showing a simple set of operations typically performed in this type of network for a 2-dimensional input.

product of the kernel and the overlapped (k, l) subsection of the input matrix in order to produce a new data matrix of size $(n - k + 1, m - l + 1)$. In other words, given the input matrix X and kernel matrix K , we can calculate the elements of the output matrix Y as:

$$Y_j^i = \sum_{a=0}^k \sum_{b=0}^l X_{(j+b)}^{(i+a)} K_b^a. \quad (4.4)$$

We see that while in a fully connected network every node in a layer affects every node in the next layer, a node in a CNN will only have an influence on, at most, $k \times l$ nodes in the convoluted layer. This is an example of *sparse connectivity*. Furthermore, it is common to subject each element of the output of the convolution to an activation function, similar to what we saw in equation (4.1), so that

$$Y_j^i = g \left(\sum_{a=0}^k \sum_{b=0}^l X_{(j+b)}^{(i+a)} K_b^a \right). \quad (4.5)$$

Another common operation in CNNs is *pooling*, which works as an attempt to summarize further the result of the convolution. Pooling functions behave much like the convolution operation, but instead of matrix multiplication we apply a correlated operation like taking

the average or maximum value of each region. This allows for some level of shift invariance, that is, the pooled layer should not change significantly if we introduce small shifts in the spatial position of the data.

There are further changes one can make to these operations, such as using a strided convolution, "skipping" some number of steps rather than applying the convolution to every possible region or padding the data before convolution so that output size is not necessarily smaller than the input size. In any case, the stereotypical base structure of the CNN tends to be stacked blocks of convolution, activation and pooling layers, such that the combination of these three operations is commonly referred to as a *convolutional layer*. These convolutional layers perform the task of feature extraction, that is, they attempt to extract the more relevant parts of some given data. It is common, however, to add some "dense" fully connected layers before the output to aid with the intended task for the model (classification, regression, etc.). These models then end up having a basic structure resembling the mock-up shown in Figure 4.3, though real applications will often have a much larger number of layers, and different variations of CNNs can have fairly distinct structures.

4.1.2 Residual Networks

Following the trend of being inspired by the human brain, ResNets try to emulate behaviour from neocortical neurons by allowing *skip connections*, where the information from a layer can skip some number of the following layers and recombine [64]. That is, if a number of layers in what we call a *residual block* acts on an input \mathbf{x} to produce the output \mathbf{y} so that

$$\mathbf{y} = f(\mathbf{x}), \tag{4.6}$$

then in a residual block (pictured in Figure 4.4 we have

$$\mathbf{y} = f(\mathbf{x}) + \mathbf{x}. \tag{4.7}$$

This means the network learns by adjusting the residual values from the input (hence the name "residual network"). Crucially, this makes the network relatively resistant to vanishing gradients, a problem that occurs when a network's weights essentially stop updating for deeper layers due to the chain-rule nature of the calculation of the derivatives. By allowing a previous layer l to have an impact on the updating of the weights of a layer $l + s$, they act to prevent an exponential diminishing of the gradient of the loss function. Neural networks are essentially universal function approximators and, in theory, an increase in number of layers

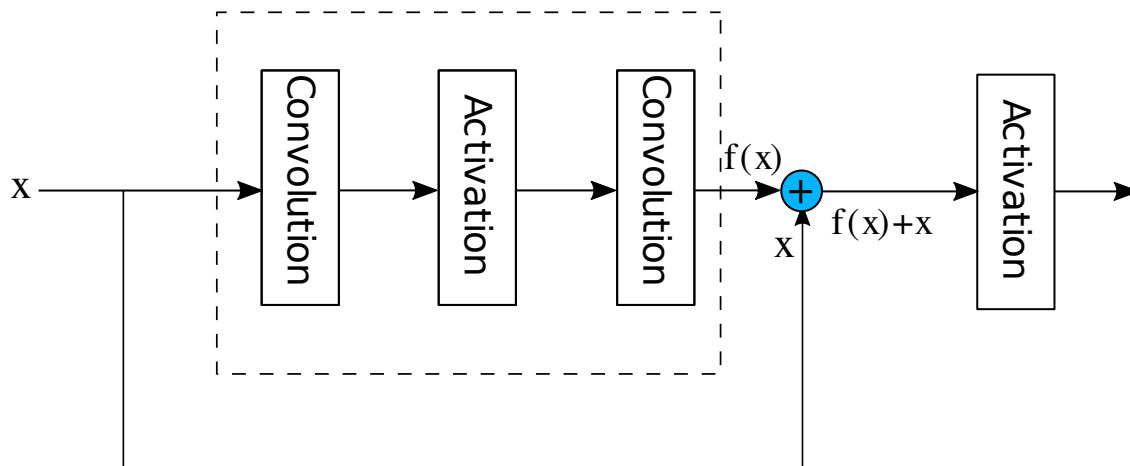


Figure 4.4: Example of a residual block.




Figure 4.5: An example of different attempts to fit a function to data. **a)** A linear function is used to approximate the data. **b)** A quadratic polynomial is used. **c)** A high-degree polynomial is used. The black circles represent the training data, while the red square represents some datum in the test set. Figure taken from [65].

should result in better accuracy: if some network architecture is capable of approximating a class of functions \mathcal{C} , then the adding of extra layers should result in an architecture capable of approximating function class \mathcal{C}' such that $\mathcal{C} \subseteq \mathcal{C}'$. However, due to propagation effects such as the vanishing gradient problem, many architectures lose performance after a certain depth. Residual networks are particularly robust against this, in part because it is easier for the optimization process to learn $f(\mathbf{x}) = 0$ such that $\mathbf{y} = f(\mathbf{x}) + \mathbf{x} = \mathbf{x}$ (this can be done by simply setting the weights of the layers to zero) than to directly learn the mapping $f(\mathbf{x}) = \mathbf{x}$. This means that, in theory, residual networks should more easily be able to approximate a larger class of functions when compared to simpler architectures.

4.2 Regularization

In addition to obtaining results on training data, machine learning algorithms need to be robust enough to be able to reliably analyze new data. To this end, there is a plethora of techniques that aim to reduce the generalization error (the error obtained when introducing new data to the algorithm), even if it means increasing the training error. This is said to reduce *overfitting*, which is to say it avoids the input→output mapping function being too specific to that particular set of training data. A common way of introducing this concept compares linear regression with polynomial regression, as shown in Figure 4.5: while the high degree polynomial provides the lowest discrepancy with the training data, its error on the test datum is as large as the linear fit. On the other hand, despite its larger error on the training set, the quadratic polynomial provides the smallest error at test time. With a naïve selection of models, the algorithm may tend to prefer the case that simply minimizes training error, but there are strategies that can be implemented in order to combat this. These strategies are collectively known as regularization, and their development is one of the major research foci in machine learning. Here we will cover two of these strategies: weight decay and dropout.

4.2.1 Weight decay

Weight decay, sometimes also known as L^2 regularization, is a regularization strategy that focuses on penalizing large weights vectors. This is motivated by the basic intuition that the simplest function possible is $f(x) = 0$, and as such the farther some given function is from zero, the more “complex” it is. Therefore, taking a simple (and not necessarily singular) interpretation, if we want to avoid excessively complex functions in our algorithm we should be able to accomplish this, at least in part, by restricting the norm of the weights vector $|\mathbf{w}|$. A simple way to do this is by adding a quadratic contribution to the loss function, as follows:

$$\mathcal{L}(\mathbf{w}, X, Y) = \mathcal{L}_0(\mathbf{w}, X, Y) + \frac{\lambda}{2} \mathbf{w}^2, \quad (4.8)$$

where λ is a chosen numerical parameter and \mathcal{L}_0 is the unregularized loss function. The $\frac{1}{2}$ factor is introduced by convention, so that when we calculate the gradient we have

$$\nabla \mathcal{L} = \nabla \mathcal{L}_0 + \lambda \mathbf{w}. \quad (4.9)$$

This way the updating of the weights vector is given by

$$\begin{aligned}\mathbf{w}' &= \mathbf{w} - \alpha \nabla \mathcal{L} \\ &= (1 - \alpha \lambda) \mathbf{w} - \alpha \nabla \mathcal{L}_0,\end{aligned}\tag{4.10}$$

and thus we see that, as long as $0 < \alpha \lambda < 1$, large components of the weights vector suffer a penalty proportional to their size during the update. This restriction can be more or less severe depending on the choice of λ , which is an experimental problem and highly network dependent.

4.2.2 Dropout

In the previous section we argued that the complexity of some function could be given in some way by its norm. However, as mentioned, this is not the only way to interpret complexity. In particular we shall now probe the issue of the complexity of a function by looking at its smoothness, and we can explore this by analyzing its robustness against small perturbations. Bishop showed in 1995 that, indeed, training on input data with artificially added noise is equivalent to training with (a generalization of) weight decay [66], yet perturbations in the training data are not the only way to implement this: we can instead introduce perturbations in a neural network itself, by randomly deactivating (or dropping out) a fraction of the network nodes when applying the algorithm to a given input during training. This technique, called *dropout* in literature, was developed in 2014 and has since then become commonplace in the design of neural network architectures [67]. Note that dropout is equivalent to applying a Bernoulli distribution with probability p and random variable r over the set of weights in a given layer:

$$\begin{aligned}w'_i &= w_i * B(p, r) \\ &= \begin{cases} w_i & \text{if } r > p, \\ 0 & \text{otherwise.} \end{cases}\end{aligned}\tag{4.11}$$

5 Towards Bayesian Neural Networks

In regular supervised learning implementations of neural networks, the output given by the network is a single value (be it some category in the case of classification problems or some set of real numbers when performing regression). It's worth taking a moment to reflect on what this means: let's imagine some network is built and trained so that it can match some recorded bird call to a certain species, and then we use it to analyse a recording of a loved one singing in the shower. Now, it might be difficult to explain, and certainly rather rude, to tell the person that according to our model they sound like a particularly screechy and annoying sort of parrot, yet without any further information that is a possible outcome, as our model does not know how to deal with human voices. Since humanity is prone to silly conflicts [68, 69], we would like to avoid any tension whenever possible. A possible way around this would be to find some strategy to quantify the uncertainty of our model, so as to clarify the fact that the network does not really know what it's looking at. The typical way of dealing with such a problem is by using Bayesian inference methods. In this section we will quickly explore Bayesian statistics and figure out a way of approximating this process using neural networks. The main references here will be [70], [71] and [72].

5.1 Bayesian modelling and the Gaussian Process

Bayesian statistics has at its kernel Bayes' theorem. For an abstract set S , we can define probability as a real-valued function that satisfies the Kolmogorov axioms:

- $\forall \sigma \subseteq S, P(\sigma) \geq 0$,
- $\forall \sigma_1, \sigma_2 \subseteq S : \sigma_1 \cap \sigma_2 = \emptyset \Rightarrow P(\sigma_1 \cup \sigma_2) = P(\sigma_1) + P(\sigma_2)$,
- $P(S) = 1$.

Then Bayes' theorem is given by

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)},$$

where

$$P(B) = \sum_i P(B|A_i)P(A_i)$$

turns out to be just a normalization constant. In the Bayesian approach this theorem is used with A standing for some hypothesis or set of parameters and B standing for some collected

data. Therefore:

$$P(\text{hypothesis}|\text{data}) \propto P(\text{data}|\text{hypothesis}) P(\text{hypothesis}) . \quad (5.1)$$

Conventionally, $P(\text{hypothesis}|\text{data})$ is known as the *posterior probability*, $P(\text{data}|\text{hypothesis})$ is the *likelihood function* and $P(\text{hypothesis})$ is the *prior probability*. This last one is not necessarily unique for a given problem: it can be seen as the degree of belief in some outcome and thus can be influenced by subjective factors. Equation (5.1) can then be seen as an update to a prior belief when faced with new data.

Let us now consider the set of N Q -dimensional inputs $X \in \mathbb{R}^{N \times Q}$ and the corresponding set of D -dimensional outputs $Y \in \mathbb{R}^{N \times D}$. We would like to find some function $f(\mathbf{x})$ in the space of functions F such that:

$$\forall \mathbf{x} \in X \exists \mathbf{y} \in Y : \mathbf{y} = f(\mathbf{x}) .$$

To accomplish this, we first consider some prior distribution $p(f)$ over F , giving us the preliminary beliefs as to which functions may be more likely to fit the data. The posterior will then be:

$$p(f|X, Y) \propto p(Y|X, f) p(f) .$$

The choice of the prior here is rather important: if we do not make any particular assumption, then any function consistent with the data will be taken as equally valid.

One way to deal with this problem would be to choose a prior that restricts the space of functions significantly (for example, demanding that $f(\mathbf{x})$ is linear). A second way of approaching the problem would be to give a prior probability to every single function in F , giving higher probabilities to less complex (e.g. smoother) functions. The issue with the first case is that, to obtain good results, we must have a good idea of how to restrict our function space from the start, as if we narrow down our space to, say, linear functions, but our data happens to be actually generated by some trigonometric relation, then predictions for new data will be inaccurate. On the other hand, if we do not narrow down the function space enough, then we will again run into the original issue of having too many candidates. The second approach seems to have a computability problem, as without restrictions we may have an infinite number of functions to compute. However, there is a way around this problem that can allow us to perform tractable computation: the Gaussian process (GP).

A Gaussian process is a generalization of the Gaussian distribution. While a distribution describes random scalar or vectorial variables, a process governs the properties of functions.

In a loose sense, we can look at a function $f(x)$ as an infinite vector consisting of the values of the function evaluated at every x . The property of GPs that will make our problem computable is the fact that, when we attempt to figure out the properties of a function at a finite number of points, inference through a Gaussian process will give an equivalent answer as if we were working with the full, infinite set of points.

In order to implement this, we want to place a joint Gaussian distribution over all function values:

$$p(Y|X) \approx \mathcal{N}(0, K(X, X)) = \sqrt{\frac{1}{2\pi K(X, X)}},$$

where $K(X, X)$ is a chosen *kernel* or *covariance function*, which defines the scalar similarity between every pair of input options⁴, effectively being equivalent to a choice of basis functions, and is given by an $N \times N$ matrix. However, the evaluation of the normal distribution implies inverting this very matrix, an operation with $\mathcal{O}(N^3)$ time complexity. This is still rather computationally intensive, so further simplification is needed.

5.2 Variational Inference and MC dropout

A way to approximate the Gaussian process described above is to restrict our functions to depend only on some set of variables ω . The distribution of predictions for some new input \mathbf{x}^* , given the initial datasets X and Y , is then

$$p(\mathbf{y}^*|\mathbf{x}^*, X, Y) = \int p(\mathbf{y}^*|\mathbf{x}^*, \omega) p(\omega|X, Y) d\omega.$$

Since $p(\omega|X, Y)$ is not analytically defined, we substitute it with some variational distribution $q(\omega)$ such that we obtain an approximate result

$$q(\mathbf{y}^*|\mathbf{x}^*) = \int p(\mathbf{y}^*|\mathbf{x}^*, \omega) q(\omega) d\omega.$$

To validate the above approximation we need to make sure that $q(\omega)$ and $p(\omega|X, Y)$ are as close as possible. This can be done by minimising the relative entropy (also known as

⁴For a more extensive discussion, chapters 2 and 4 of [72] are recommended.

Kullback–Leibler (KL) divergence)

$$\text{KL}(q(\omega)||p(\omega | X, Y)) = \sum_{\omega} q(\omega) \log \left(\frac{q(\omega)}{p(\omega | X, Y)} \right),$$

which is a measurement of the difference between both distributions. This has been shown to be equivalent to maximizing the *log evidence lower bound*,

$$\mathcal{L}_{VI} = \int q(\omega)p(Y | X, \omega)d\omega - \text{KL}(q(\omega)||p(\omega)),$$

a process that is known as variational inference. It was shown by Gal and Ghahramani in 2016 that with the appropriate choice of covariance function, the Gaussian process can be approximated such that the log evidence lower bound is proportional to

$$\mathcal{L}_{VI} \propto -\frac{1}{2N} \sum_{n=1}^N \|f(x_n) - y_n\|_2^2 - \frac{p_1}{2\tau N} \|\mathbf{M}_1\|_2^2 - \frac{p_2}{2\tau N} \|\mathbf{M}_2\|_2^2,$$

and that the parameter matrices $\mathbf{M}_1^{\text{opt}}$ and $\mathbf{M}_2^{\text{opt}}$ that maximize this quantity will be equivalent to the parameters that minimize the mean squared error (MSE) loss for a NN with a single hidden layer making use of dropout and weight decay,

$$\mathcal{L}_D = \frac{1}{2N} \sum_{n=1}^N \|f(x_n) - y_n\|_2^2 + \lambda_1 \|\mathbf{W}_1\|_2^2 + \lambda_2 \|\mathbf{W}_2\|_2^2,$$

where \mathbf{W}_1 and \mathbf{W}_2 stand for the weight matrices in between layers [73]. While this result is, as mentioned, valid for NNs with a single hidden layer, it is extendable to arbitrary depth and even convolutional architectures [74]. This shows that, while typically only used during training, keeping a proper implementation of dropout active during inference and passing the same inputs forward a large number of times will yield a distribution that approximates the Gaussian process method, giving a natural way to deduce the uncertainty of a network for some result (at the cost of increasing evaluation time).




Figure 6.1: Example of a spectrogram produced for a simulation of the coalescence of a binary system with $\mathcal{M} = 20 M_{\odot}$. Note the similarity with Figure 2.1.

6 Methods

The intent of the work developed in the context of this dissertation was to explore new machine learning methods for treating gravitational wave data from ground-based GW interferometers such as the advanced LIGO and Virgo detectors. Specifically, we will study the application of mainstream computer vision (CV) tools such as the residual networks discussed in 4.1.2 to the time-frequency representation of gravitational waves (“spectrograms”, such as seen in Figure 6.1). These spectrograms are 2-dimensional data and can be produced in such a way as to be suitable for interpretation by CV techniques: the majority of digital images have a colour depth of 8 bits, meaning that for some image with width W , height H and colour channels N (typically there are 3 colour channels: red, green and blue. Some images have an additional channel, called the alpha channel, which controls transparency), its mathematical representation is a $W \times H \times N$ tensor whose components are 8-bit integers, that is, integers in the $[0, 255]$ range. The generation of a spectrogram from some given timeseries generates a $W_s \times H_s$ tensor (the dimensions can be adjusted depending on the desired time/frequency resolution and ranges), and although the resulting component values (which translate the power registered at some time and for a certain frequency) are floating

numbers, they can be normalized to the 8-bit integers. While this does mean we will lose some resolution in the power values, normalizing in this way makes this sort of data an innate fit for tools used in computer vision. In particular, since a lot of existing machinery is built to deal with RGB images, we have a natural way of incorporating data from multiple detectors by stacking independent spectrograms into a single image. This is explained further in 6.1.2.

6.1 Generating datasets

We begin by describing the generation of the datasets used in our analysis. Generation of GW data through solving the Einstein equations directly (numerical relativity (NR) waveforms) is very computationally expensive and is not feasible on a local consumer-level machine. When large amounts of GW data is needed, it is common to generate events using some relatively easy to evaluate function, called an approximant, that outputs a waveform with significant overlap with the NR result. In our case, all CBC waveforms employed in the classification datasets were obtained using pyCBC [75] with the SEOBNRv4_ROM approximant while the regression datasets use SEOBNRv4HM_ROM [76], IMRPhenomPv2 [77] and IMRPhenomD [78]. For the sake of simplicity we start by considering spinless black holes and quasi-circular binaries with no orbital eccentricity. Furthermore, since current GW detector networks are far more sensitive to the plus polarization than to the cross one, we only generate plus-polarized waves. Harking back to equation (2.24), this has the drawback of making impossible to break the degeneracy between luminosity distance and inclination.

6.1.1 Single detector waveforms for classification

The purpose of this first dataset is to allow our DL models to discern the presence of a GW signal with data collected from a single detector. In particular we employ a 500 s noise segment from the Hanford detector with initial GPS time $t_{\text{GPS}} = 1187058342$ s. Defining τ to be some time from the start of our noise segment, we randomly select $\tau_0 \in [5, 495]$ s and isolate the window $[\tau_0 - 5 \text{ s}, \tau_0 + 5 \text{ s}]$. This strain selection, which we denote as n , is then whitened through inverse spectrum truncation, using its own amplitude spectral density (ASD) (that is, the signal is fourier transformed to the frequency domain, divided by the ASD, then transformed back to time domain). Then, we apply a bandpass filter from 20 Hz to 300 Hz, as well as notch filters at the individual frequencies 60 Hz, 120 Hz and 240 Hz. For the generation of the waveform signal strain h , a random pair of black hole masses $(m_1, m_2) \in U([5, 100]) M_{\odot}$ is selected for a BBH merger with luminosity distance $d_L = 2000$ Mpc and inclination $\iota = \frac{\pi}{2}$. This waveform is whitened using the ASD of the

selected noise strain n and the same filtering process is undertaken. The resulting waveform is injected into the noise window in such a way that the maximum amplitude occurs at τ_0 . Following this, the constant-Q transform [79] is calculated for the $[\tau_0 - 0.16 \text{ s}, \tau_0 + 0.4 \text{ s}]$ interval in the composite signal $S = h + t$, and a spectrogram is produced. A second spectrogram without signal injection is also generated for the same interval. Both spectrograms are saved as images and appropriately labeled as “signal” and “background”. This process is iterated 5000 times to build our dataset. The same procedure is taken for the luminosity distances $d_L = 100, 300, 1000, 1500$ and 2000 Mpc. A summary of the single detector classification procedure is shown in the first row of Table 6.1.

6.1.2 Multiple detector waveforms for classification

In order to combine the data from all three detectors (Hanford, Livingston and Virgo) we select coincident segments of 500 s from all detectors starting at a certain t_{GPS} time. The process is then identical to that of a single detector case with τ_0 , m_1 and m_2 randomly generated. A $[\tau_0 - 5 \text{ s}, \tau_0 + 5 \text{ s}]$ time window is extracted from the longer segments for the three detectors. The resulting background strain data, n_H , n_L and n_V , for Hanford, Livingston and Virgo interferometers respectively, is treated in the same way as described above. However, when injecting a signal, one must make sure that the specific ASD of each detector is being used. After the generation of the signal waveform and its injection into the background noise segments, we include the antenna power from each detector into our time series S_H , S_L and S_V . At this stage, we can emulate the sky position for the signal by randomly choosing one of three detectors as a reference, and shift the beginning of the other two time series according to their time delay with respect to the reference detector. Once the three spectrograms are produced they are combined into a $560 \times 560 \times 3$ array in such a way that each of them is represented by a certain colour channel in a RGB image. Specifically, Hanford, Livingston and Virgo datasets are mapped into the Red, Green and Blue channels respectively, as can be seen in Figure 6.2(a). As in the single detector case, an equivalent background spectrogram without signal injection is produced. Both, background (left) and signal (right) spectrograms, are represented in Figure 6.2(b). Once again, this process is iterated 5000 times for each luminosity distance we consider. The classification procedure for the multiple detector case is summarized in the second row of Table 6.1.




Figure 6.2: **(a)** Combining single detector spectrogram data into a single RGB image, to be used by the deep learning networks. The Hanford (top-left), Livingston (top-right) and Virgo (bottom-left) spectrogram data, are used as the Red, Green and Blue images, respectively, to build the full RGB image (bottom-right). **(b)** RGB image from background ('bg') labeled spectrogram (left) as compared with a spectrogram labeled as signal ('sig') from a GW injection into noise (right).

6.1.3 Mass dependence dataset

The purpose of this dataset is to check how the trained models perform depending on the BH binary’s component masses and luminosity distance, both for the single and multiple detector cases. To this end, we consider a number of mass combinations, where for each m_1 , ranging from 5 to 100 M_\odot in steps of 2 M_\odot , there is a set of values for m_2 , in the range $[m_1, 100] M_\odot$, covered also with steps of 2 M_\odot . Regarding the luminosity distance d_L , we use distances ranging from 100 to 2000 Mpc, in steps of 100 Mpc. The inclination is kept fixed at $\iota = \frac{\pi}{2}$. For each d_L and each of the 1225 (m_1, m_2) mass combinations, a waveform is generated and injected into the detector’s noise following the same procedure as described above. However, in this case, only the spectrograms with the injections are saved and labeled with the corresponding distance and mass.

6.1.4 Regression datasets

We want to understand if DL methods can be used to extract information about the physical parameters from the generated spectrograms. This procedure is typically denoted as regression. For this purpose, a larger dataset was deemed necessary and only the multiple detector case, was considered. To avoid a dependence on a particular approximant, three different datasets for each of the approximants `SEOBNRv4HM_ROM`, `IMRPhenomPv2` and `IMRPhenomD` were built. It is also relevant to note that, due to an apparent degradation of the regression close to the upper range of the sampled distances, we decided to consider distances up to 4 Gpc, although we set our range of operation to go up to a maximum distance of 2.5 Gpc. Here, we do not build separate datasets for particular values of the distance but instead let d_L be randomly generated within this range. The component masses m_1 and m_2 are again randomly sampled in the $[5, 100] M_\odot$ interval while the inclination takes a random value in the $[0, \pi]$ interval. We also sample the sky position by taking into account the antenna pattern of each detector. In our regression dataset black holes are assumed to have a dimensionless spin in the range $[-1, 1]$ and those are aligned with the orbital angular momentum, allowing the computation of the effective inspiral spin, χ_{eff} ,

$$\chi_{\text{eff}} \approx \frac{m_1 \chi_1^\parallel + m_2 \chi_2^\parallel}{m_1 + m_2}, \quad (6.1)$$

where χ_i^\parallel is the component of the i -th spin along the orbital angular momentum. Since we assume from the beginning that a given input to the regression model will necessarily contain a GW signal of some sort, we need not to worry about generating the background-only cases.

Classification		
Parameters	Train size	Validation size
Single detector $(m_1, m_2) \sim U(5, 100) M_\odot$, $d_L = [100, 300, 1000, 1500, 2000]$ Mpc, $\iota = \frac{\pi}{2}$, approximant: SEOBNRv4_ROM 560×560 pixels 8-bit gray scale	4000 images	1000 images
Total images	20000	5000
Multiple detector (same parameters as above) 560×560 pixels 8-bit RGB	4000 images	1000 images
Total images	20000	5000

Table 6.1: Description of the classification datasets for training and validation with both single and multiple detectors. The images are generated from the wave-forms calculated by pyCBC. For the classification datasets, the individual masses (m_1, m_2) are sampled with a uniform distribution within the range of 5 to 100 M_\odot .

Furthermore we impose a threshold for the signal-to-noise ratio (SNR) so that we only allow cases where $\text{SNR} > 5$. The sizes of the SEOBNRv4HM_ROM, IMRPhenomPv2 and IMRPhenomD datasets are, respectively, 15961, 12922, and 13281. An extra dataset focused on lower masses ($m_1, m_2 \in [5, 35]M_\odot$) was also generated using the SEOBNRv4HM_ROM approximant, containing 15538 events. This was combined with the original SEOBNRv4HM_ROM dataset for a lower-mass weighted dataset with a total of 31499 items. All this information is summarized in Table 6.1.

6.2 Building a network for classification

6.2.1 Single-detector training distance comparison

For the task of classification, we chose to use a residual network architecture as described in 4.1.2. We use the `fastai` library for python, which allows for the quick building of a DL network and gives access to some pre-trained ResNet architectures. Using the single-detector classification datasets described in 6.1, we set up a series of 34-layer networks, trained for 10 epochs using cyclical learning rates [80] between 2×10^{-3} and 2×10^{-1} and using L2 regularization with $\lambda = 1 \times 10^{-5}$, in order to explore how the signal injection distance affects network performance on injections at different distances. This is portrayed in Figure 6.3,

Regression		
Parameters	Train size	Validation size
Multiple detector $(m_1, m_2) \sim U(5, 100) M_\odot$, $d_L \sim U(100, 4000) \text{ Mpc}$, $\iota \sim U(0, \pi)$, $\chi_{\text{eff}} \sim U(-1, 1)$, restriction: $\text{SNR} > 5$ approximant: <code>SEOBNRv4HM_ROM</code> 224×224 pixels 8-bit RGB	12769 images	3192 images
(same parameters as above) approximant: <code>IMRPhenomPv2</code>	10338 images	2584 images
(same parameters as above) approximant: <code>IMRPhenomD</code>	10625 images	2656 images
Total images	43009	14689
Extra dataset		
Multiple detector $(m_1, m_2) \sim U(5, 35) M_\odot$, $d_L \sim U(100, 4000) \text{ Mpc}$, $\iota \sim U(0, \pi)$, $\chi_{\text{eff}} \sim U(-1, 1)$, restriction: $\text{SNR} > 5$ approximant: <code>SEOBNRv4HM_ROM</code> 224×224 pixels 8-bit RGB	15538 images	

Table 6.2: Description of the regression datasets for training and validation. The images are generated from the wave-forms calculated by pyCBC. For the regression dataset the parameters of individual masses, distances, inclination and spin are uniformly sampled. The extra dataset is generated in order to include more examples of small masses and complement the dataset generated with the `SEOBNRv4HM_ROM` approximant.




Figure 6.3: Performance of the different classifier networks. **a)** The accuracy of the differently trained networks on events injected at the same distance as each NN’s training distance. **b)** Accuracy of the networks when tested on events injected at 2 Gpc. **c)** Accuracy of the network trained with 2 Gpc injections on events injected at different distances.

where accuracy refers to a metric calculated as

$$\text{Acc} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}},$$

where TP (TN) corresponds to the correctly identified signals (backgrounds) and FP (FN) correspond to incorrectly classified signals (backgrounds). We found that, while networks trained on closer (louder) injections showed very high accuracy (over 99%) for loud events, they were not able to detect signal injections at larger distances, showing a dismal 50% accuracy on 2 Gpc injections, no better than a coin toss. On the other hand we found that while training on more distant (quieter) injections yielded a lower accuracy than the first case, at 72.2%, this network retained the ability to perform its function competently over all tested injection distances. Following the previously mentioned principle that a residual network should always retain or increase its performance when more layers are added, we finally choose to train a 101-layer ResNet on the 2 Gpc classification dataset. Though this did not yield an increase in performance regarding the accuracy metric, the training time is not prohibitively long, and as such we use this architecture moving forward.

6.2.2 Comparing single-detector and multiple-detector performance

Having settled on an architecture and a training set, we now want to compare the classification performance when using data from a single detector in a grayscale image versus using data




Figure 6.4: Simulated signal scores using one single detector for different luminosity distances and evaluated with DL networks trained with GW waveforms from BBH mergers at a luminosity distance of 2000 Mpc. Results are shown as a function of the BH masses of the binary system, m_1 and m_2 , for GW signals from sources at 400 Mpc (left), 1000 Mpc (middle) and 2000 Mpc (right).




Figure 6.5: Simulated signal scores using multiple detectors for different luminosity distances and evaluated with DL networks trained with GW waveforms from BBH mergers at a luminosity distance of 2000 Mpc. Results are shown as a function of the BH masses of the binary system, m_1 and m_2 , for GW signals from sources at 400 Mpc (left), 1000 Mpc (middle) and 2000 Mpc (right).

from multiple detectors in an RGB image. Training a 101-layer ResNet for 5 epochs with RGB spectrograms of signal injections at 2 Gpc for 10 epochs, we obtain an accuracy of 82.4%, a 14% performance increase when compared to the equivalent single-detector case. This is further illustrated when we test both the single-detector and multiple-detector networks on the datasets generated in 6.1.3, as shown in Figures 6.4 and 6.5: for injections at some given distance, not only is the range of detectable mass combinations (i.e. score > 0.5) larger, the scores for the multiple detector network are overall higher than the single detector case, implying this network is generally more confident than its predecessor. Hence we choose to use multiple detector data for our work.

Having settled on an architecture for our classifier, it is now useful to look at the receiver




Figure 6.6: ROC curve for the best-performing classifier on multiple-detector 2000 Mpc data. The red star displays the current threshold location on the ROC used for classifying the events into signal (score ≥ 0.5) or background (score ≤ 0.5).

operating characteristic (ROC) curve of the network in Figure 6.6, regarding the 2000 Mpc data. We can see that using the current score threshold for detection we obtain a signal efficiency $\epsilon_S = \frac{TP}{TP+FP} = 72\%$ as well as a background rejection $1 - \epsilon_B = \frac{FP}{FP+TN} = 92\%$. However, by adjusting the threshold, we could optimize the classifier’s performance to our need, depending on whether we would like to be more or less strict with the detection process.

6.3 Building a network for parameter estimation

For the task of parameter estimation (or regression) we will test the ability of our model to retrieve physical quantities from spectrogram data. Our target parameters for this work are the luminosity distance d_L , the chirp mass of the binary system \mathcal{M} , the NAP captured by the detector, and the effective inspiral spin χ_{eff} . We again choose to use a residual-based architecture, although there are a few key differences when compared to the classifier network:

- We use a cross-residual network (xResNet), a tweaked version of ResNets that can enable greater network generalization [81].

- We use an 18-layer network, as increasing depth in residual networks can actually hamper regression performance [82].
- We use blur pooling instead of simple max pooling layers, to make the network more resistant to translational shifts in input data [83].
- We place dropout layers with probability 0.5 after convolutions and before pooling, as described in [74], in order to obtain an approximately Bayesian network.

For each of the regression datasets described in 6.1.4, a network was trained for 7 epochs using a cyclical learning rate with an initial value of 1×10^{-2} and using 1×10^{-3} as the λ parameter for weight decay. Due to the usage of blur pooling and the change in architecture, this network is more computationally expensive to train than the classifier network. To compensate for this, we downsample our input spectrograms to a resolution of 128×128 .

Classification		
Base architecture	Hyperparameters	Accuracy
ResNet-101 + custom header	input size: $275 \times 275 \times 3(1)$, batch size: 8 images, learning rate: $[2 \times 10^{-3}, 2 \times 10^{-1}]$, weight decay: 1×10^{-5} Cross Entropy Loss (CE)	Single: 0.72 Multiple: 0.82
Regression		
Base architecture	Hyperparameters	
xResNet-18 + Blur average layer + MC Dropout + custom header	input size: $128 \times 128 \times 3$, batch size: 64 images, learning rate: 1×10^{-2} , weight decay: 1×10^{-3} , Mean Squared Error loss (MSE)	

Table 6.3: Convolutional Neural Networks architectures employed for the classification, multiple (single) detectors, and regression tasks. The first two columns show the setup of the networks, while the last column in the classification table shows the achieved accuracy. The custom header for the classification CNN is described in (TODO), the custom header for the regression model has the same structure with the main difference being that the final layer has only one unit with a linear activation function.

6.3.1 Parameter estimation tests on synthetic injections

Figures 6.7 to 6.10 show the networks' behaviour on injected signals from each network's validation set, for varying values of signal to noise ratio (SNR). To take advantage of MC




Figure 6.7: Calibration results for d_L using the different approximant datasets, SEOBNRv4HM_ROM (top), IMRPhenomPv2 (middle) and IMRPhenomD (bottom), and for different SNR thresholds.




Figure 6.8: Same as 6.7 but showing the calibration results for the network antenna power.




Figure 6.9: Same as 6.7 but for the chirp mass in the source frame.




Figure 6.10: Same as 6.7 but for the effective inspiral spin.




Figure 6.11: Mean absolute error (MAE) of the d_L predictions as a function of SNR, for the SEOBNRv4HM_ROM dataset.

dropout, each spectrogram in the dataset was forward-passed through the network 100 times, with the dropout layers stochastically turning on and off at each pass. For each event, the predictions represented in the figures are the mean value of the 100 evaluations.

As a general rule, we do not notice significant differences between networks, suggesting that, at least in this parameter space, the choice of approximant is not a particularly impactful factor. As such, Figures 6.11 to 6.14 show the mean absolute error (MAE) for the predictions as a function of SNR only for the SEOBNRv4HM_ROM dataset.

For the case of the luminosity distance, we see that for the lower SNR threshold we obtain a significant prediction spread for larger injection distances. As we increase the SNR threshold we start obtaining distribution closer to the ideal behaviour (represented by the dashed white line). However, note that due to the $\frac{1}{d_L}$ dependence in the amplitude, the requirement of a high SNR implies, on average, a restriction to lower distances. Figure 6.11 shows that the MAE of the predictions starts at around 0.4 Gpc for the lowest SNRs and tends to the order of 0.1 Gpc for higher SNRs.

The distribution of predictions for the NAP shows a similar dependence on SNR, however it is noticeable that with the lowest SNR threshold, along with the distribution of events around the diagonal line, there is a population of predictions around the horizontal line at $\text{NAP} \approx 0.6$. This may be a sign of a fundamental limitation: for lower SNR values, the likelihood of an event only being clearly visible in a single detector is higher, and in such cases we would not expect there to be enough information to be able to reconstruct the NAP. Since this population cannot be discerned in Figure 6.12, interpretations of the MAE for $\text{SNR} \in [5, 10]$ are difficult.

The predictions for the chirp mass in the source frame (Figure 6.9) turn out to be quite




Figure 6.12: Mean absolute error (MAE) of the NAP predictions as a function of SNR, for the SEOBNRv4HM_ROM dataset.




Figure 6.13: Mean absolute error (MAE) of the \mathcal{M} predictions as a function of SNR, for the SEOBNRv4HM_ROM dataset.




Figure 6.14: Mean absolute error (MAE) of the χ_{eff} predictions as a function of SNR, for the SEOBNRv4HM_ROM dataset.

impressive even for lower SNR thresholds, with a very tight prediction distribution, which becomes even more accurate as SNR increases. Analysis of Figure 6.13 reveals that the worst case scenario for the MAE of the prediction is a deviation of $4 M_{\odot}$, decreasing to around $2.5 M_{\odot}$ for higher SNRs. This is a remarkable result.

The predictions for the effective inspiral spin (Figure 6.10) follow the behaviour seen for the chirp mass, though with a larger spread. In Figure 6.14 we see once again a significant drop of the MAE of the prediction as SNR increases. It is worth noting that this is a higher order effect in the GW waveform, and the ability to detect these subtler effects is in itself worthy of highlight.

7 Results and discussion

Finally, we shall test our classifier and regression networks on real GW data released by the LIGO/Virgo collaboration. We noticed at this point that the network was hesitant to assign low chirp mass values to the events, and on further investigation realized that low masses were quite under-represented in our datasets (likely due to the requirement that $\text{SNR} > 5$). This was the reason the low-mass extra dataset in Table 6.1 was produced, and a new network was trained on a superset composed of the original `SEOBNRv4HM_ROM` data concatenated with the extra dataset providing the low-mass correction, using a 70/30 train/validation split. Initially, this analysis encompassed only the ten confident BBH events in GWTC-1 as well as the “marginal confidence” from the same catalog [84], comprised by the O1 and O2 runs. However, on the 27th of October, 2020, the LVC released a new catalog, GWTC-2, containing thirty-seven new BBH events [85]. Despite having had our networks trained on O2 conditions (the O3 run has a significantly reduced noise floor), we managed to, in a matter of hours from the release, do a full analysis of the new catalog, obtaining satisfactory results as will be described below. While ideally the networks should be trained on data under the same conditions, the fact that these tools can be so quickly applied and show already some degree of generalization is highly encouraging.

7.1 Classification network

To analyse real GW events we follow a procedure similar to what was described in 6.1, fetching the strain data for the available detectors at the GPS time described in the released data and combining the individual spectrograms into RGB images. Since we trained our detector on BBH signals, we do not analyse data known to belong to binary neutron star or neutron star-black hole events, leaving out GW170817 from GWTC-1 as well as GW190425 and GW190426 from GWTC-2. GW190814 is included, although one of the objects involved is officially estimated to have a mass around $2.6 M_{\odot}$, which is generally deemed too heavy for a neutron star but too light for a black hole [86].

GWTC-1 Confident		GWTC-1 Marginal		GWTC-2			
Event	Score	Event	Score	Event	Score	Event	Score
GW170814	1.00	MC151116	0.73	GW190521	1.00	GW190708_232457	0.98
GW150914	1.00	MC161217	0.72	GW190602_175927	1.00	GW190909_114149	0.97
GW170823	1.00	MC170705	0.51	GW190424_180648	1.00	GW190514_065416	0.96
GW170104	1.00	MC170630	0.49	GW190620_030421	1.00	GW190814	0.95
GW170729	0.99	MC170219	0.45	GW190503_185404	1.00	GW190521_074359	0.95
GW170809	0.97	MC161202	0.40	GW190727_060333	1.00	GW190731_140936	0.92
GW151012	0.96	MC170423	0.35	GW190929_012149	1.00	GW190513_205428	0.92
GW170608	0.92	MC170208	0.33	GW190915_235702	1.00	GW190421_213856	0.87
GW170818	0.88	MC170720	0.30	GW190630_185205	1.00	GW190412	0.81
GW151226	0.87	MC151012A	0.26	GW190519_153544	1.00	GW190728_064510	0.77
-	-	MC151008	0.20	GW190706_222641	1.00	GW190719_215514	0.76
-	-	MC170405	0.14	GW190413_134308	1.00	GW190803_022701	0.66
-	-	MC170616	0.12	GW190701_203306	1.00	GW190930_133541	0.58
-	-	MC170412	0.09	GW190517_055101	1.00	GW190828_065509	0.56
-	-	-	-	GW190408_181802	1.00	GW190924_021846	0.40
-	-	-	-	GW190910_112807	1.00	GW190707_093326	0.35
-	-	-	-	GW190828_063405	0.99	GW190720_000836	0.16
-	-	-	-	GW190413_052954	0.99	-	-
-	-	-	-	GW190512_180714	0.98	-	-
-	-	-	-	GW190527_092055	0.98	-	-

Table 7.1: Classifier scores for GWTC-1 marginal detections (left), GWTC-1 confident detections (middle) and GWTC-2 detections (right). Scores that pass the classifier’s threshold for detection (score ≥ 0.5) are in bold.

The results from our classifier are displayed in Table 7.1. All the confident BBH detections from the LVC for GWTC-1 are corroborated by the classifier network, with GW170814, GW150914, GW170823 and GW170104 getting the maximum score of 1.00, while GW151226 is given a score of 0.87, the lowest of the lot. For the GWTC-1 marginal events we get lower scores across the board and indeed most of the events are rejected, though there are three events (MC151116, MC161217 and MC170705) which pass the threshold for detection and

may warrant further analysis. The results for the GWTC-2 catalog are also impressive, especially taking into account the lack of optimization: of the thirty-seven events analysed, thirty-four are corroborated by the network, twenty-seven of which are given scores above the 0.9 mark. Over both catalogs, this amounts to a $\sim 94\%$ accuracy for our classifier.

As an aside remark, it is interesting to note that we have also obtained high scores for signals proposed in alternative GW catalogs [87, 88]. For example, our method yields a score of 0.75 for GW151216, proposed in [88].


7.2 Regression network

7.2.1 GWTC-1


To perform parameter estimation on the GWTC-1 confident detections, we must forward-pass the spectrograms of the events through a model several times to make use of the Bayesian properties of MC dropout. The mean and standard deviation of the predictions can then be calculated. For our analysis we chose to do a total of 1500 passes per event to be safe, though this is probably excessive and we have obtained similar results using only 200 passes. Our initial results are presented in Figure 7.1, with the circular marker representing the mean, the blue box representing the 90% interval, and the red error bars delimiting the 3σ range⁵. This is plotted against the 90% confidence interval (in green) reported in the LVC data release. The top panel shows predictions obtained by passing each event 500 times to each of the 3 networks trained on different approximants, then combining the resulting distributions. The bottom panel shows the predictions using the network trained on the low-mass-corrected dataset, with each BBH event spectrogram forward-passed through the network 1500 times, again making use of MC dropout to obtain a distribution of outputs.

Chirp mass For the case of the chirp mass we find that, using the 3 network combination, prediction distributions for 6 of the 10 GWTC-1 events are incompatible with the published LVC data, as represented on the leftmost panel of Figure 7.1a. This is particularly noticeable in the cases of GW151226 and GW170608: the network is unable to correctly predict lower chirp-mass values. We expect this to be related to the lack of lower mass representation in the original datasets. In fact, when we analyse the same data with the network trained on

⁵It is important to note that we do not expect the uncertainty we obtain from MC dropout to be the full uncertainty associated with some prediction. Crucially, MC dropout assumes homoskedasticity, that is, equal variance (or noise) throughout all data, which is not necessarily our case. Furthermore, there may be sources of variance that MC dropout does not take into account.



(a)



(b)

Figure 7.1: Predictions of the DL network for the chirp mass (left), luminosity distance (middle) and effective inspiral spin (left). The top panels (a) do not include the low-mass correction while the bottom panels (b) do.

the low-mass corrected dataset, we see a considerable improvement. Looking at the leftmost panel of Figure 7.1b, we now find that our MC dropout 3σ range for the chirp mass predictions overlaps with the LVC 90% confidence interval for all 10 GWTC-1 events.

Luminosity distance The middle panel of Figure 7.1a displays the combined results for the three approximants without the low-mass distribution, for the d_L of the GWTC-1 confident BBH detections. For the luminosity distance most of our predictions are compatible with the LVC predictions up to a network uncertainty of 3σ , with the single exception being GW170823. However, when we use network trained with the low-mass correction, as displayed in the middle panel of Figure 7.1b, all of our predictions become compatible with the LVC values.

Effective inspiral spin When inferring the effective inspiral spin χ_{eff} using the combined results for the three approximants and without including the low-mass distribution, we find six events with a significant disagreement with published results (in the same sense as discussed previously for the chirp mass), as can be seen in the rightmost panel of Figure 7.1a. As shown in the corresponding panel of Figure 7.1b, when the network is trained with the addition of the low-mass dataset, our results improve. All events, except for GW151226, show compatibility between the LVC 90% confidence interval and the MC dropout 3σ uncertainty.

7.2.2 GWTC-2

Finally we will discuss the application of our best-performing network (the one trained on the low-mass corrected dataset) to the BBH events from the GWTC-2 catalog, stressing the point that our network was not optimized for the noise conditions of the O3 run. As previously done, we perform parameter estimation on the chirp mass, luminosity distance and effective inspiral spin, forward-passing the spectrograms of each event through the network 1500 times. The results and comparison with LVC values are plotted in Figure 7.2. For the chirp mass, we find our MC dropout 3σ range to be compatible with the LVC 90% confidence interval for 33 of the 37 analysed events. We find a very decent performance in estimating the effective inspiral spin as well, with 29 compatible events. Our worst performance in the GWTC-2 events comes in the estimation of the luminosity distance, with only 24 compatible events. Since all of the incompatibilities come from underestimations of the LVC values, this may be related to the fact that the distance’s effect on the waveform is simply as a scalar overall factor. As the noise floor in O3 is lower than in O2 and O1, whose data was used



Figure 7.2: Best-performing DL network's prediction for the chirp mass (left), luminosity distance (middle) and effective inspiral spin (left) for the BBH events from GWTC-2.

for training, and given the fact that our method for producing the spectrograms casts the power amplitudes to the 8-bit integers in a relative way, the network may be interpreting a signal with a higher SNR due to a smaller noise factor as a signal from a smaller distance. Nonetheless, the fact that the network can infer physical parameters such as chirp mass and spin with some accuracy despite no optimization for the noise conditions reinforces the potential inherent to these methods.

8 Conclusions

The goal of this dissertation was to develop a novel manner of performing detection and parameter inference of BBH events. While current methods for detection [89, 90, 91, 92] and parameter estimation [93, 94] exist, the work presented here can serve as the basis of a pipeline that is able to easily integrate classification and regression tasks. We achieve this by studying spectrograms (save of signals in the time-frequency domain, and training neural networks on synthetic signal injections to be able to perform the required tasks.

We started our investigations by training a series of ResNets on signal injections into real detector strain data at different distances. We found that training a network for far-away, fainter signals maintained a good level of performance even for closer, louder signals, while the opposite was not true. Furthermore, we noticed that combining the spectrograms of all 3 available detectors into a single RGB image and training our network on such data provided a significant boost to the network’s accuracy. On application to real BBH events from the LVC’s GWTC-1 catalog, the network corroborated all the confident detections and identified 3 candidates from the marginal confidence set as GW signals. Applied to GWTC-2 data without optimization for different noise conditions, 34 out of 37 candidates for BBH events pass the network’s threshold for detection.

For the task of parameter inference, or regression, we trained an 18-layer xResNet on RGB spectrograms of signal injections with varying physical spins, mass, distance and sky position, setting distance, chirp mass, network antenna power (as a proxy for sky position) and effective inspiral spin as our regression targets. Furthermore, we use MC dropout to obtain a natural estimation of the network’s uncertainty. At a fledgling level of development, we show that it is possible to retrieve the relevant parameters from the data with varying levels of accuracy. Though we found particular success in the case of the chirp mass, it is worth noting that the effective inspiral spin only has influence in higher-order terms of the GW waveform, and so the ability to resolve it highlights the presented method’s sensitivity. Performing parameter inference on GWTC-1, we found remarkable agreement with published data: All of our prediction distributions for the chirp mass and luminosity distance are compatible with the LVC’s 90% confidence interval up to an MC dropout uncertainty of 3σ . Among the effective inspiral spin predictions, only the distribution for GW151226 is incompatible with the published data. For GWTC-2, again without optimization for different noise conditions, from a total of 37 BBH events we find 33, 29 and 24 compatible events for our predictions on the chirp mass, effective inspiral spin and distance, respectively.

It is important to note that despite the initial promising results, there is a lot of room for

future work:

- While deep residual networks were used for the work presented here, a more in-depth architecture search could provide a model with better performance.
- By storing our spectrograms as RGB images, we limit ourselves to the 8-bit integer range. Furthermore, the normalization used (where for each detector's spectrogram the minimum value is cast to 0 and the maximum value is cast to 255) means we lose track of the absolute value of the signal power. This also means that the level of noise can be suppressed in high SNR injections, which may undermine MC dropout performance as it assumes homoskedastic noise. Adapting the networks to use pure floating-point spectrogram tensors, without normalization, may allow for an increase in performance. This is work in progress at this point.
- This type of classifier can in principle be extended to look at different morphologies of signal (such as neutron star mergers or core-collapse supernovae). On the other hand, it could also be used to detect and classify the different types of glitches that plague GW interferometers. Some initial explorations of this have shown promise.

References

- [1] Emília Pásztor and Curt Roslund. “An interpretation of the Nebra disc”. en. In: *Antiquity* 81.312 (June 2007). Publisher: Cambridge University Press, pp. 267–278. ISSN: 0003-598X, 1745-1744. DOI: 10.1017/S0003598X00095168. URL: <https://www.cambridge.org/core/journals/antiquity/article/abs/an-interpretation-of-the-nebra-disc/3FFD318A44C610CFC69D327FD12C530E> (visited on 04/27/2021).
- [2] T. P. Ray. “The winter solstice phenomenon at Newgrange, Ireland: accident or design?” en. In: *Nature* 337.6205 (Jan. 1989). Number: 6205 Publisher: Nature Publishing Group, pp. 343–345. ISSN: 1476-4687. DOI: 10.1038/337343a0. URL: <https://www.nature.com/articles/337343a0> (visited on 04/27/2021).
- [3] O. Neugebauer. “Studies in Ancient Astronomy. VIII. The Water Clock in Babylonian Astronomy”. en. In: *Isis* (Oct. 2015). Publisher: History of Science Society. ISSN: 0021-1753. DOI: 10.1086/347965. URL: <https://www.journals.uchicago.edu/doi/abs/10.1086/347965> (visited on 04/27/2021).
- [4] Teije de Jong. “A study of Babylonian planetary theory I. The outer planets”. en. In: *Arch. Hist. Exact Sci.* 73.1 (Jan. 2019), pp. 1–37. ISSN: 1432-0657. DOI: 10.1007/s00407-018-0216-0. URL: <https://doi.org/10.1007/s00407-018-0216-0> (visited on 04/27/2021).
- [5] David Pingree. “The Teaching of the Almagest in Late Antiquity”. In: *Apeiron* 27.4 (Jan. 1994). ISSN: 2156-7093, 0003-6390. DOI: 10.1515/APEIRON.1994.27.4.75. URL: <https://www.degruyter.com/document/doi/10.1515/APEIRON.1994.27.4.75/html> (visited on 04/27/2021).
- [6] George Saliba. “The Role of the Almagest Commentaries in Medieval Arabic Astronomy: A Preliminary Survey of Ṭūsī’s Redaction of Ptolemy’s Almagest”. In: *A History of Arabic Astronomy: Planetary Theories During the Golden Age of Islam* (1995). ISBN: 0814738893 Publisher: NYU Press.
- [7] René Taton, Curtis Wilson, and Michael Hoskin. *Planetary Astronomy from the Renaissance to the Rise of Astrophysics, Part A, Tycho Brahe to Newton*. Cambridge University Press, 2003. ISBN: 0-521-54205-7.
- [8] Urbain Le Verrier. “Lettre de M. Le Verrier à M. Faye sur la théorie de Mercure et sur le mouvement du périhélie de cette planète”. In: *Comptes rendus hebdomadaires des séances de l’Académie des sciences (Paris)* 49 (1859), pp. 379–383.

- [9] A. Einstein. “Die Grundlage der allgemeinen Relativitätstheorie”. en. In: *Annalen der Physik* 354.7 (1916). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.19163540702>, pp. 769–822. ISSN: 1521-3889. DOI: <https://doi.org/10.1002/andp.19163540702>. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/andp.19163540702> (visited on 04/27/2021).
- [10] Slava G. Turyshev. “Experimental Tests of General Relativity”. In: *Annual Review of Nuclear and Particle Science* 58.1 (2008). Publisher: Annual Reviews, pp. 207–248. ISSN: 0163-8998. DOI: 10.1146/annurev.nucl.58.020807.111839. URL: <https://dx.doi.org/10.1146/annurev.nucl.58.020807.111839>.
- [11] John Michell. “VII. On the means of discovering the distance, magnitude, &c. of the fixed stars, in consequence of the diminution of the velocity of their light, in case such a diminution should be found to take place in any of them, and such other data should be procured from observations, as would be farther necessary for that purpose. By the Rev. John Michell, B.D. F.R.S. In a letter to Henry Cavendish, Esq. F.R.S. and A.S”. In: *Philosophical Transactions of the Royal Society of London* 74 (Jan. 1784). Publisher: Royal Society, pp. 35–57. DOI: 10.1098/rstl.1784.0008. URL: <https://royalsocietypublishing.org/doi/10.1098/rstl.1784.0008> (visited on 04/27/2021).
- [12] Pierre Simon Laplace. “Beweis des Satzes, dass die anziehende Kraft bey einem Weltkörper so groß seyn könne, dass das Licht davon nicht ausströmen kann”. In: *Allgemeine Geographische Ephemeriden, Vol. 4, Issue 1, p. 1-6* 4 (July 1799), pp. 1–6. URL: <http://adsabs.harvard.edu/abs/1799A11GE...4...1L> (visited on 05/12/2021).
- [13] Colin Montgomery, Wayne Orchiston, and Ian Whittingham. “Michell, Laplace and the origin of the black hole concept”. en. In: *Journal of Astronomical History and Heritage* 12.2 (2009), pp. 90–96.
- [14] K. Schwarzschild. “On the Gravitational Field of a Mass Point According to Einstein’s Theory”. In: *Abh. Konigl. Preuss. Akad. Wissenschaften Jahre 1906,92, Berlin,1907* 1916 (1916). URL: <http://adsabs.harvard.edu/abs/1916AbhKP1916...189S> (visited on 04/27/2021).
- [15] S. Antoci and D.-E. Liebscher. “Editor’s Note: On the Gravitational Field of a Mass Point According to Einstein’s Theory by K. Schwarzschild”. en. In: *General Relativity and Gravitation* 35.5 (May 2003), pp. 945–950. ISSN: 1572-9532. DOI: 10.1023/A:

1022919909683. URL: <https://doi.org/10.1023/A:1022919909683> (visited on 04/27/2021).
- [16] S. Antoci. “David Hilbert and the origin of the "Schwarzschild solution"”. In: *Meteorological and Geophysical Fluid Dynamics* (2004), p. 343. URL: <http://adsabs.harvard.edu/abs/2004mgfd.book..343A> (visited on 04/27/2021).
- [17] Tarun Biswas. “Physical Interpretation of Coordinates for the Schwarzschild Metric”. In: *arXiv pre-print server* (2008). URL: <https://arxiv.org/abs/0809.1452>.
- [18] S. W. Hawking. “The Nature of Space and Time”. In: *arXiv:hep-th/9409195* (Sept. 1994). arXiv: hep-th/9409195. URL: <http://arxiv.org/abs/hep-th/9409195> (visited on 04/27/2021).
- [19] Roger Penrose. “Gravitational Collapse and Space-Time Singularities”. In: *Phys. Rev. Lett.* 14.3 (Jan. 1965). Publisher: American Physical Society, pp. 57–59. DOI: 10.1103/PhysRevLett.14.57. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.14.57> (visited on 04/27/2021).
- [20] Nobel Media AB. *Press Release: The Nobel Prize in Physics 2020*. en-US. URL: <https://www.nobelprize.org/prizes/physics/2020/press-release/> (visited on 01/08/2021).
- [21] Albert Einstein. “Naherungsweise Integration der Feldgleichungen der Gravitation”. In: *Sitzungsberichte der Koniglich Preuischen Akademie der Wissenschaften (Berlin), Seite 688-696*. (1916). URL: <http://adsabs.harvard.edu/abs/1916SPAW.....688E> (visited on 04/27/2021).
- [22] Arthur Stanley Eddington. “The propagation of gravitational waves”. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 102.716 (Dec. 1922). Publisher: Royal Society, pp. 268–282. DOI: 10.1098/rspa.1922.0085. URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1922.0085> (visited on 04/27/2021).
- [23] Albert Einstein, Max Born, and Hedwig Born. *The Born-Einstein letters: friendship, politics, and physics in uncertain times : correspondence between Albert Einstein and Max and Hedwig Born from 1916 to 1955 with commentaries by Max Born*. Translated from the German. pp. 121 OCLC: 225215993. Basingstoke, Hampshire; New York: Macmillan, 2005. ISBN: 978-1-4039-4496-2 978-1-349-72911-1.

- [24] F. A. E. Pirani. “On the Physical significance of the Riemann tensor”. In: *Acta Physica Polonica* 15 (1956), pp. 389–405. URL: <http://adsabs.harvard.edu/abs/1956AcPP..15..389P> (visited on 04/27/2021).
- [25] Conference on the Role of Gravitation in Physics et al., eds. *The role of gravitation in physics: report from the 1957 Chapel Hill Conference*. English. OCLC: 1226071248. 2017. ISBN: 978-3-945561-29-4.
- [26] R. A. Hulse and J. H. Taylor. “Discovery of a pulsar in a binary system”. en. In: *ApJ* 195 (Jan. 1975), p. L51. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/181708. URL: <http://adsabs.harvard.edu/doi/10.1086/181708> (visited on 01/08/2021).
- [27] Joel M. Weisberg, Joseph H. Taylor, and Lee A. Fowler. “Gravitational Waves from an Orbiting Pulsar”. In: *Sci Am* 245.4 (Oct. 1981), pp. 74–82. ISSN: 0036-8733. DOI: 10.1038/scientificamerican1081-74. URL: <http://www.nature.com/doi/10.1038/scientificamerican1081-74> (visited on 01/08/2021).
- [28] Nobel Media AB. *Press Release: The Nobel Prize in Physics 1993*. en-US. URL: <https://www.nobelprize.org/prizes/physics/1993/press-release/> (visited on 01/08/2021).
- [29] Chiang-Mei Chen, James M. Nester, and Wei-Tou Ni. “A brief history of gravitational wave research”. In: *Chinese Journal of Physics* 55.1 (2017). Publisher: Elsevier BV, pp. 142–169. ISSN: 0577-9073. DOI: 10.1016/j.cjph.2016.10.014. URL: <https://dx.doi.org/10.1016/j.cjph.2016.10.014>.
- [30] The LIGO Scientific Collaboration. “Advanced LIGO”. In: *Class. Quantum Grav.* 32.7 (Apr. 2015). arXiv: 1411.4547, p. 074001. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/32/7/074001. URL: <http://arxiv.org/abs/1411.4547> (visited on 02/25/2021).
- [31] F Acernese et al. “Advanced Virgo: a second-generation interferometric gravitational wave detector”. In: *Class. Quantum Grav.* 32.2 (Jan. 2015), p. 024001. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/32/2/024001. URL: <https://iopscience.iop.org/article/10.1088/0264-9381/32/2/024001> (visited on 02/25/2021).
- [32] Richard Brito et al. “Proca stars: Gravitating Bose–Einstein condensates of massive spin 1 particles”. en. In: *Physics Letters B* 752 (Jan. 2016), pp. 291–295. ISSN: 03702693. DOI: 10.1016/j.physletb.2015.11.051. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0370269315009077> (visited on 01/04/2021).

- [33] P. C. Peters and J. Mathews. “Gravitational Radiation from Point Masses in a Keplerian Orbit”. In: *Physical Review* 131 (July 1963), pp. 435–440. ISSN: 1536-6065. DOI: 10.1103/PhysRev.131.435. URL: <http://adsabs.harvard.edu/abs/1963PhRv..131..435P> (visited on 05/18/2021).
- [34] V. Pierro and I. M. Pinto. “Exact solution of Peters-Mathews equations for any orbital eccentricity”. In: *Il Nuovo Cimento B Series 11* 111.5 (1996). Publisher: Springer Science and Business Media LLC, pp. 631–644. ISSN: 1826-9877. DOI: 10.1007/bf02726655. URL: <https://dx.doi.org/10.1007/bf02726655>.
- [35] Chad Hanna et al. “Method to estimate ISCO and ring-down frequencies in binary systems and consequences for gravitational wave data analysis”. In: *Class. Quantum Grav.* 26.1 (Jan. 2009). arXiv: 0801.4297, p. 015009. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/26/1/015009. URL: <http://arxiv.org/abs/0801.4297> (visited on 05/31/2021).
- [36] Luc Blanchet. “Gravitational Radiation from Post-Newtonian Sources and Inspiralling Compact Binaries”. en. In: *Living Rev. Relativ.* 17.1 (Dec. 2014), p. 2. ISSN: 1433-8351. DOI: 10.12942/lrr-2014-2. URL: <https://doi.org/10.12942/lrr-2014-2> (visited on 05/28/2021).
- [37] Dayton C. Miller. “The Ether-Drift Experiment and the Determination of the Absolute Motion of the Earth”. en. In: *Rev. Mod. Phys.* 5.3 (July 1933), pp. 203–242. ISSN: 0034-6861. DOI: 10.1103/RevModPhys.5.203. URL: <https://link.aps.org/doi/10.1103/RevModPhys.5.203> (visited on 01/07/2021).
- [38] Michele Maggiore. *Gravitational Waves*. Oxford University Press, Oct. 2007. ISBN: 978-0-19-857074-5. DOI: 10.1093/acprof:oso/9780198570745.001.0001. URL: <https://oxford.universitypressscholarship.com/view/10.1093/acprof:oso/9780198570745.001.0001/acprof-9780198570745> (visited on 01/03/2021).
- [39] Bernard F. Schutz. “Networks of gravitational wave detectors and three figures of merit”. In: *Class. Quantum Grav.* 28.12 (June 2011). arXiv: 1102.5421, p. 125023. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/28/12/125023. URL: <http://arxiv.org/abs/1102.5421> (visited on 01/11/2021).
- [40] Piotr Jaranowski, Andrzej Królak, and Bernard F. Schutz. “Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection”. en. In: *Phys. Rev. D* 58.6 (Aug. 1998), p. 063001. ISSN: 0556-2821, 1089-4918. DOI: 10.1103/

- PhysRevD.58.063001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.58.063001> (visited on 01/12/2021).
- [41] Lee Samuel Finn. “Aperture synthesis for gravitational-wave data analysis: Deterministic sources”. en. In: *Phys. Rev. D* 63.10 (Apr. 2001), p. 102001. ISSN: 0556-2821, 1089-4918. DOI: 10.1103/PhysRevD.63.102001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.63.102001> (visited on 01/12/2021).
- [42] Stephen Fairhurst. “Localization of transient gravitational wave sources: beyond triangulation”. In: *Class. Quantum Grav.* 35.10 (May 2018), p. 105002. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/1361-6382/aab675. URL: <https://iopscience.iop.org/article/10.1088/1361-6382/aab675> (visited on 01/12/2021).
- [43] The Virgo Collaboration. *Advanced Virgo Technical Design Report*. Tech. rep. VIR-0128A-12. Apr. 2012. URL: <https://tds.virgo-gw.eu/ql/?c=8940>.
- [44] F Matichard. “Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance”. en. In: (2015), p. 31.
- [45] F. Acernese et al. “Measurements of Superattenuator seismic isolation by Virgo interferometer”. en. In: *Astroparticle Physics* 33.3 (Apr. 2010), pp. 182–189. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2010.01.006. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650510000253> (visited on 02/28/2021).
- [46] Jennifer C. Driggers, Jan Harms, and Rana X. Adhikari. “Subtraction of Newtonian noise using optimized sensor arrays”. en. In: *Phys. Rev. D* 86.10 (Nov. 2012), p. 102001. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.86.102001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.86.102001> (visited on 02/25/2021).
- [47] Scott A. Hughes and Kip S. Thorne. “Seismic gravity-gradient noise in interferometric gravitational-wave detectors”. en. In: *Phys. Rev. D* 58.12 (Nov. 1998), p. 122002. ISSN: 0556-2821, 1089-4918. DOI: 10.1103/PhysRevD.58.122002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.58.122002> (visited on 02/25/2021).
- [48] D. V. Martynov et al. “The Sensitivity of the Advanced LIGO Detectors at the Beginning of Gravitational Wave Astronomy”. In: *Phys. Rev. D* 93.11 (June 2016). arXiv: 1604.00439, p. 112004. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.93.112004. URL: <http://arxiv.org/abs/1604.00439> (visited on 02/25/2021).

- [49] Carlton M. Caves. “Quantum-mechanical noise in an interferometer”. en. In: *Phys. Rev. D* 23.8 (Apr. 1981), pp. 1693–1708. ISSN: 0556-2821. DOI: 10.1103/PhysRevD.23.1693. URL: <https://link.aps.org/doi/10.1103/PhysRevD.23.1693> (visited on 02/26/2021).
- [50] Thomas Corbitt and Nergis Mavalvala. “Quantum noise in gravitational-wave interferometers: Overview and recent developments”. en. In: *arXiv:gr-qc/0306055* (May 2003). arXiv: gr-qc/0306055, p. 23. DOI: 10.1117/12.507487. URL: <http://arxiv.org/abs/gr-qc/0306055> (visited on 02/25/2021).
- [51] P. Fritschel et al. *Instrument Science White Paper 2020*. Tech. rep. LIGO-T2000407. Aug. 2020. URL: <https://dcc.ligo.org/LIGO-T2000407/public>.
- [52] Gabriela González. “Suspensions Thermal Noise in the LIGO Gravitational Wave Detector”. en. In: *Class. Quantum Grav.* 17.21 (Nov. 2000). arXiv: gr-qc/0006053, pp. 4409–4435. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/17/21/305. URL: <http://arxiv.org/abs/gr-qc/0006053> (visited on 05/31/2021).
- [53] M Lorenzini and the Virgo Collaboration. “The monolithic suspension for the Virgo interferometer”. en. In: *Class. Quantum Grav.* 27.8 (Apr. 2010), p. 084021. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/27/8/084021. URL: <https://iopscience.iop.org/article/10.1088/0264-9381/27/8/084021> (visited on 05/31/2021).
- [54] Roger P. Netterfield et al. “Low mechanical loss coatings for LIGO optics: progress report”. en. In: ed. by Michael L. Fulton and Jennifer D. T. Kruschwitz. San Diego, California, USA, Aug. 2005, 58700H. DOI: 10.1117/12.618288. URL: <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.618288> (visited on 02/28/2021).
- [55] Gabriele Vajente. “Overview of coating thermal noise”. en. In: *Gravitational Wave Advanced Detector Workshop 2018* (May 2018). LIGO Document G1800920-v4, p. 19.
- [56] B. P. Abbott et al. “A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals”. en. In: *Class. Quantum Grav.* 37.5 (Feb. 2020). Publisher: IOP Publishing, p. 055002. ISSN: 0264-9381. DOI: 10.1088/1361-6382/ab685e. URL: <https://doi.org/10.1088/1361-6382/ab685e> (visited on 02/20/2021).
- [57] Alejandro Torres-Forné et al. “Application of dictionary learning to denoise LIGO’s blip noise transients”. en. In: *Phys. Rev. D* 102.2 (July 2020). arXiv: 2002.11668, p. 023011. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.023011. URL: <http://arxiv.org/abs/2002.11668> (visited on 02/28/2021).

- [58] Jade Powell et al. “Classification methods for noise transients in advanced gravitational-wave detectors II: performance tests on Advanced LIGO data”. en. In: *Class. Quantum Grav.* 34.3 (Feb. 2017), p. 034002. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/1361-6382/34/3/034002. URL: <https://iopscience.iop.org/article/10.1088/1361-6382/34/3/034002> (visited on 05/12/2021).
- [59] Miquel Llorens-Montegudo et al. “Classification of gravitational-wave glitches via dictionary learning”. en. In: *Class. Quantum Grav.* 36.7 (Apr. 2019). arXiv: 1811.03867, p. 075005. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/1361-6382/ab0657. URL: <http://arxiv.org/abs/1811.03867> (visited on 02/28/2021).
- [60] Michael Zevin et al. “Gravity Spy: Integrating Advanced LIGO Detector Characterization, Machine Learning, and Citizen Science”. In: *Class. Quantum Grav.* 34.6 (Mar. 2017). arXiv: 1611.04596, p. 064003. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/1361-6382/aa5cea. URL: <http://arxiv.org/abs/1611.04596> (visited on 03/02/2021).
- [61] Chris Smith et al. “The History of Artificial Intelligence”. en.
- [62] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. Adaptive computation and machine learning. Cambridge, Massachusetts: The MIT Press, 2016. ISBN: 978-0-262-03561-3.
- [63] Aston Zhang et al. *Dive into Deep Learning*. 2020.
- [64] Thomson. “Neocortical layer 6, a review”. In: *Front. Neuroanat.* (2010). ISSN: 16625129. DOI: 10.3389/fnana.2010.00013. URL: <http://journal.frontiersin.org/article/10.3389/fnana.2010.00013/abstract> (visited on 01/04/2021).
- [65] Benyamin Ghogh and Mark Crowley. “The Theory Behind Overfitting, Cross Validation, Regularization, Bagging, and Boosting: Tutorial”. en. In: *arXiv:1905.12787 [cs, stat]* (May 2019). arXiv: 1905.12787. URL: <http://arxiv.org/abs/1905.12787> (visited on 01/16/2021).
- [66] Chris M. Bishop. “Training with Noise is Equivalent to Tikhonov Regularization”. en. In: *Neural Computation* 7.1 (Jan. 1995), pp. 108–116. ISSN: 0899-7667, 1530-888X. DOI: 10.1162/neco.1995.7.1.108. URL: <https://www.mitpressjournals.org/doi/abs/10.1162/neco.1995.7.1.108> (visited on 01/16/2021).
- [67] Nitish Srivastava et al. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958. URL: <http://jmlr.org/papers/v15/srivastava14a.html>.

- [68] George Garnett. “ROBERT CURTHOSE: THE DUKE WHO LOST HIS TROUSERS”. In: *Anglo-Norman Studies 35*. Ed. by David Bates. Vol. 35. Proceedings of the Battle Conference 2012. Boydell & Brewer, 2013, pp. 213–244. ISBN: 978-1-84383-857-9. DOI: 10.7722/j.ctt31nhmm.17. URL: <http://www.jstor.org/stable/10.7722/j.ctt31nhmm.17>.
- [69] *Looking back: Australia’s Emu Wars*. en-AU. Oct. 2016. URL: <https://www.australiangeographic.com.au/topics/wildlife/2016/10/on-this-day-the-emu-wars-begin/> (visited on 01/17/2021).
- [70] David Barber. *Bayesian Reasoning and Machine Learning*. en. Cambridge: Cambridge University Press, 2011. ISBN: 978-0-511-80477-9. DOI: 10.1017/CB09780511804779. URL: <http://ebooks.cambridge.org/ref/id/CB09780511804779> (visited on 01/18/2021).
- [71] Yarin Gal. “Uncertainty in Deep Learning”. en. PhD thesis. University of Cambridge, Sept. 2016. URL: <http://www.cs.ox.ac.uk/people/yarin.gal/website/thesis/thesis.pdf>.
- [72] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. en. Adaptive computation and machine learning. OCLC: ocm61285753. Cambridge, Mass: MIT Press, 2006. ISBN: 978-0-262-18253-9.
- [73] Yarin Gal and Zoubin Ghahramani. “Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning”. en. In: *arXiv:1506.02142 [cs, stat]* (Oct. 2016). arXiv: 1506.02142. URL: <http://arxiv.org/abs/1506.02142> (visited on 01/20/2021).
- [74] Yarin Gal and Zoubin Ghahramani. “Bayesian Convolutional Neural Networks with Bernoulli Approximate Variational Inference”. en. In: *arXiv:1506.02158 [cs, stat]* (Jan. 2016). URL: <http://arxiv.org/abs/1506.02158> (visited on 08/16/2020).
- [75] Alex Nitz et al. *gwastro/pycbc: PyCBC release v1.16.11*. Version v1.16.11. Oct. 2020. DOI: 10.5281/zenodo.4075326. URL: <https://doi.org/10.5281/zenodo.4075326>.
- [76] Roberto Cotesta, Sylvain Marsat, and Michael Pürrer. “Frequency domain reduced order model of aligned-spin effective-one-body waveforms with higher-order modes”. In: *Phys. Rev. D* 101.12 (2020), p. 124040. DOI: 10.1103/PhysRevD.101.124040. arXiv: 2003.12079 [gr-qc].

- [77] Mark Hannam et al. “Simple Model of Complete Precessing Black-Hole-Binary Gravitational Waveforms”. In: *Phys. Rev. Lett.* 113 (15 Oct. 2014), p. 151101. DOI: 10.1103/PhysRevLett.113.151101. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.113.151101>.
- [78] Sascha Husa et al. “Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and anatomy of the signal”. In: *Phys. Rev. D* 93 (4 Feb. 2016), p. 044006. DOI: 10.1103/PhysRevD.93.044006. URL: <https://link.aps.org/doi/10.1103/PhysRevD.93.044006>.
- [79] Judith C. Brown. “Calculation of a constant Q spectral transform”. In: *The Journal of the Acoustical Society of America* 89.1 (Jan. 1991). Publisher: Acoustical Society of America, pp. 425–434. ISSN: 0001-4966. DOI: 10.1121/1.400476. URL: <https://asa.scitation.org/doi/10.1121/1.400476> (visited on 06/01/2021).
- [80] Leslie N. Smith. “No More Pesky Learning Rate Guessing Games”. In: *CoRR* abs/1506.01186 (2015). arXiv: 1506.01186. URL: <http://arxiv.org/abs/1506.01186>.
- [81] Brendan Jou and Shih-Fu Chang. “Deep Cross Residual Learning for Multitask Visual Recognition”. en. In: *arXiv:1604.01335 [cs]* (July 2016). URL: <http://arxiv.org/abs/1604.01335> (visited on 08/18/2020).
- [82] Dongwei Chen et al. “Deep Residual Learning for Nonlinear Regression”. en. In: *Entropy* 22.2 (Feb. 2020), p. 193. ISSN: 1099-4300. DOI: 10.3390/e22020193. URL: <https://www.mdpi.com/1099-4300/22/2/193> (visited on 02/20/2021).
- [83] Richard Zhang. “Making Convolutional Networks Shift-Invariant Again”. en. In: *arXiv:1904.11486 [cs]* (June 2019). URL: <http://arxiv.org/abs/1904.11486> (visited on 08/18/2020).
- [84] B. P. Abbott et al. “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs”. In: *Physical Review X* 9.3, 031040 (July 2019), p. 031040. DOI: 10.1103/PhysRevX.9.031040. arXiv: 1811.12907 [astro-ph.HE].
- [85] R. Abbott et al. “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run”. In: *arXiv e-prints*, arXiv:2010.14527 (Oct. 2020), arXiv:2010.14527. arXiv: 2010.14527 [gr-qc].
- [86] R. Abbott et al. “GW190814: Gravitational Waves from the Coalescence of a 23 Solar Mass Black Hole with a 2.6 Solar Mass Compact Object”. In: *ApJ* 896.2, L44 (June 2020), p. L44. DOI: 10.3847/2041-8213/ab960f. arXiv: 2006.12611 [astro-ph.HE].

- [87] Alexander H. Nitz et al. “1-OGC: The First Open Gravitational-wave Catalog of Binary Mergers from Analysis of Public Advanced LIGO Data”. In: *The Astrophysical Journal* 872.2 (Feb. 2019), p. 195. ISSN: 1538-4357. DOI: 10.3847/1538-4357/ab0108. URL: <http://dx.doi.org/10.3847/1538-4357/ab0108>.
- [88] Barak Zackay et al. “Highly spinning and aligned binary black hole merger in the Advanced LIGO first observing run”. In: *Physical Review D* 100.2 (July 2019). ISSN: 2470-0029. DOI: 10.1103/physrevd.100.023007. URL: <http://dx.doi.org/10.1103/PhysRevD.100.023007>.
- [89] S. A. Usman et al. “The PyCBC search for gravitational waves from compact binary coalescence”. In: *Classical and Quantum Gravity* 33.21, 215004 (Nov. 2016), p. 215004. DOI: 10.1088/0264-9381/33/21/215004. arXiv: 1508.02357 [gr-qc].
- [90] Cody Messick et al. “Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data”. en. In: *Phys. Rev. D* 95.4 (Feb. 2017), p. 042001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.95.042001. URL: <https://link.aps.org/doi/10.1103/PhysRevD.95.042001> (visited on 03/08/2021).
- [91] M. Drago et al. “Coherent WaveBurst, a pipeline for unmodeled gravitational-wave data analysis”. In: *arXiv:2006.12604 [gr-qc]* (Jan. 2021). arXiv: 2006.12604. URL: <http://arxiv.org/abs/2006.12604> (visited on 03/08/2021).
- [92] T Adams et al. “Low-latency analysis pipeline for compact binary coalescences in the advanced gravitational wave detector era”. In: *Class. Quantum Grav.* 33.17 (Sept. 2016), p. 175012. ISSN: 0264-9381, 1361-6382. DOI: 10.1088/0264-9381/33/17/175012. URL: <https://iopscience.iop.org/article/10.1088/0264-9381/33/17/175012> (visited on 03/08/2021).
- [93] Gregory Ashton et al. “BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy”. In: *ApJS* 241.2, 27 (Apr. 2019), p. 27. DOI: 10.3847/1538-4365/ab06fc. arXiv: 1811.02042 [astro-ph.IM].
- [94] J. Veitch et al. “Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library”. In: *Phys. Rev. D* 91.4, 042003 (Feb. 2015), p. 042003. DOI: 10.1103/PhysRevD.91.042003. arXiv: 1409.7215 [gr-qc].