
Gravitational lensing of gravitational waves: Theoretical review and identification with deep learning

Auteur : p210248

Promoteur(s) : Cudell, Jean-Rene

Faculté : Faculté des Sciences

Diplôme : Master en sciences spatiales, à finalité approfondie

Année académique : 2021-2022

URI/URL : <http://hdl.handle.net/2268.2/14778>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.

MASTER THESIS

UNIVERSITY OF LIÈGE - FACULTY OF SCIENCES



**Gravitational lensing of gravitational waves:
Theoretical review and identification with deep learning**

Author

OFFERMANS Arthur

For a Master's degree in Space Sciences,
Research Focus

Supervisor

CUDELL Jean-René

ACADEMIC YEAR 2021–2022

Acknowledgements

First, I would like to thank Prof. Jean-René Cudell for his availability, for his advice that has been of great help for the realisation of this work, and for giving me access to the computational resource I needed.

I am also grateful to Grégory Baltus for suggestions on how to treat the problem and his help with the use of software dedicated to gravitational waves, and to Dr. Atri Bhattacharya for the technical support.

I also thank Arnaud Delaunoy and Prof. Gilles Louppe for their advice on the design of the neural network.

Then, I would like to thank the members of the evaluation committee, Dr. Michaël De Becker, Dr. Maxime Fays and Dr. Dominique Sluse, for accepting to read this master thesis.

Finally, I am deeply thankful to my friends and my family who supported me during my studies. I would not have come this far without them.

Contents

Introduction	1
1 A few preliminary concepts	3
1.1 Brief introduction to general relativity	3
1.2 Cosmological considerations	6
1.2.1 Cosmological models	6
1.2.2 Cosmological redshift	8
1.2.3 Distances in cosmology	9
1.3 Summary of conventions and important relations	10
2 Electromagnetic gravitational lensing	11
2.1 Gravitational lensing effect	11
2.2 Lens equation	12
2.3 Time delays	15
2.4 Deflection angle	17
2.5 Magnification and distortion	19
2.6 Types and number of images	20
2.7 Point-mass lens	22
3 Gravitational lensing of gravitational waves: theoretical approach	25
3.1 Introduction to gravitational waves	25
3.1.1 Gravitational waves	25
3.1.2 Wave equation	26
3.1.3 Solution for the Minkowski metric	27
3.2 Differences with the electromagnetic case	29
3.3 Taking wave effects into account	30
3.4 Wave equation with a point-mass lens	33
3.5 Analytical solution for a point-mass lens	35
3.5.1 Derivation of the solution	35
3.5.2 Re-expressing the solution	38
3.5.3 Interpreting the solution	41
3.5.4 Validity and comments	43

4	Gravitational lensing of gravitational waves: practical approach	45
4.1	Introduction to deep learning	45
4.1.1	Supervised learning	45
4.1.2	Deep learning	48
4.1.3	Convolutional Neural Networks (CNN)	51
4.2	Gravitational-wave observations	53
4.3	Importance of artificial intelligence in gravitational wave detection	55
4.3.1	Detection of gravitational waves: matched filtering	55
4.3.2	Deep learning: filtering and early warning	57
4.4	Identification of lensed gravitational waves	58
4.5	A neural network to find them	59
4.5.1	Formulation of the problem	59
4.5.2	Data generation	60
4.5.3	Model and training	63
4.5.4	Results	65
4.5.5	Discussion	69
5	Interests and future prospects	73
5.1	Advantages of the identification of lensed waves	73
5.2	Problems with the misidentification of lensed waves	74
5.3	Designing a better network	75
	Conclusion	77
	Appendix A Derivation of relations in electromagnetic lensing	79
A.1	Time delay and lensing potential	79
A.2	Deflection angle for a point-mass lens	83
A.3	Magnification for a point-mass lens	84
	Appendix B Derivation of the wave equation	85
	Appendix C Rewriting the wave equation	91
	Bibliography	93

Introduction

Gravitational waves have long been a mere prediction of general relativity as an oscillation of spacetime that propagates through it, until their indirect detection from a pulsar binary by Hulse and Taylor [1] in 1974 that was thought to lose energy through gravitational radiation [2]. One will have to wait until 2015 to see the first direct detection of a gravitational wave: GW150914 [3]. Since then, 90 waves have been detected [4] using laser interferometers. Nowadays, there are 5 such detectors: LIGO Livingston and Hanford [5, 6] (USA), Virgo [7, 8] (Italy), KAGRA [9] (Japan) and GEO 600 [10] (Germany).

Now that we are convinced of the existence of these waves and that we can detect them, some of their predicted properties and behaviours can be verified. In particular, it is predicted that gravitational waves, just like light, can be deflected by a massive body. Such a phenomenon was indeed observed with light in 1979 by Walsh, Carswell and Weymann [11]. This is called gravitational lensing, when one or several possibly deformed images of the same source reach the observer because of the deflection by a massive object. Hence, according to general relativity, such a phenomenon should also be possible with gravitational waves. A question then arises: are some of the detected waves images of the same source? This question has not found an answer yet and the search for such *lensed waves* is still going on [12]. There are, in fact, many other questions. What would a lensed wave look like? What would be the consequence of detecting or not lensed waves? What information could we retrieve from such events? How can we identify them? Why should we care? These are the questions that we will try to answer in this work, or at least that we will discuss.

The identification of gravitational lensing of gravitational waves is motivated by promising scientific prospects, among which we find testing general relativity and measuring the Hubble constant. The road that leads to it is however full of challenges. One of them is the noise, which is the main limitation to the sensitivity of the detectors. Another one is the increasing number of events, and thus of possible groups of images that may come from the same source. If one only considers pairs of waves, the number of them to analyse goes as the square of the number of events. Improvements in current detectors and the appearance of new generation detectors, such as the Einstein Telescope [13] and Cosmic Explorer [14], will increase the number of events significantly. The Einstein Telescope is indeed expected to have a rate of about $10^5 - 10^6$ binary black hole coalescence events per year [15], which corresponds to about 10^9-10^{10} to $10^{11}-10^{12}$ pairs! This highlights the need for fast algorithms that can, for example, reduce the number of pairs and keep only the promising ones. Such algorithms do not need to be optimal [13], given the purpose they are designed for.

Machine learning may be a solution to this problem. In particular, we shall investigate

deep learning. Such methods were already considered in gravitational-wave data analysis in different contexts and for different tasks [16]. It thus seems natural to evaluate its performance in the identification of lensed waves.

The purpose of this thesis is to give an overview and basic understanding of the phenomenon of gravitational lensing applied to gravitational waves, and to propose a concrete way to identify it in the detections. It must be noted that this suggestion is merely a proof of concept and that further studies will be required to design a proper solution. This work thus requires basic knowledge in both theoretical and practical aspects of very different fields, which each reader may not be familiar with. Therefore, a non-negligible part is dedicated to the explanation of basic concepts of these fields. Similarly, quite a few mathematical details may be found all along the thesis. It is however not required to read each and every line of them to understand the results and the underlying concepts. These details are yet kept for those who would be interested in knowing where all the results come from. Some developments can also be found in the appendices, because they are either too heavy and cumbersome, or simply not necessary for the purpose of this work.

In the first chapter, we shall review a few theoretical concepts that are required to understand the developments that are presented in the work, in particular general relativity and cosmology. The conventions used in the work are also introduced in this chapter.

After that, in Chapter 2, the lensing effect applied to light will be reviewed, for its results also apply to the lensing of gravitational waves. A basic understanding of this phenomenon will thus be helpful when the case of gravitational waves will be explored. We will explain lensing both conceptually and mathematically.

Then, Chapter 3 will be dedicated to the theory of lensing of gravitational waves. We shall first introduce these waves and then explore if their lensing can be different from the electromagnetic case. The wave equation accounting for the presence of a massive object will be derived and solved. The result, describing the lensing phenomenon, will then be discussed.

Once we have understood the bases of the theoretical concepts, we shall pursue our analysis of the lensing effect with its practical aspects in Chapter 4, *i.e.* how it can be identified in the observations of gravitational waves. A preliminary introduction to some concepts of machine and deep learning is required before diving into the core of the subject. We will also review how gravitational waves are detected, with or without machine learning, in order to understand better both the challenges of identifying lensed waves and the importance of techniques such as deep learning. Finally, we shall have a brief overview of current techniques to solve the initial task, and a neural network designed in the context of this work will be presented and discussed, along with its results.

The last chapter will highlight the importance and interest of finding lensed waves. We shall also discuss future prospects and how the proposed neural network can be improved.

Chapter 1

A few preliminary concepts

In this chapter, I will review some basic concepts that will be useful to understand the content of this work. I will only write about physical concepts that are not the main subjects of the thesis. Other concepts, such as gravitational waves and machine learning will be introduced properly in their dedicated sections.

1.1 Brief introduction to general relativity

General relativity is the currently accepted theory of gravitation that successfully and accurately describes phenomena that were not (or not accurately) predicted by the Newtonian theory. We will not see in detail what this theory is, but we need a few key ideas to understand and describe the gravitational lensing effect. A few conventions (*e.g.* the definition of the Minkowski metric, Riemann tensor, etc.) used in this work will also be presented. A summary of the conventions is shown in Section 1.3. This section is based on [17], where more information can be found, as well as in [18].

In this theory, it is postulated that space and time are dimensions of the same spacetime. One can represent an *event* in this spacetime by a point denoted by the coordinates

$$x^\mu = (x^0, x^1, x^2, x^3), \quad (1.1)$$

where $x^0 = ct$ is the time coordinate, with c the speed of light, and x^i the spatial coordinates (greek indices range from 0 to N , while roman ones range from 1 to N , with $N + 1$ the dimension of the spacetime). A common convention is to use natural units, *i.e.* $c = 1$, which is what will be used in this work. We then write $x^0 = t$. A spacetime is described by its *metric* $g_{\mu\nu}$, which defines distances in this spacetime. This tensor allows one to generalise the scalar product as

$$\vec{x} \cdot \vec{y} = \sum_i \sum_j g_{ij} x^i y^j \equiv g_{ij} x^i y^j \equiv x_i y^i. \quad (1.2)$$

The norm induced by this product is then

$$\|\vec{x}\| = \sqrt{\vec{x} \cdot \vec{x}} = \sqrt{g_{ij} x^i x^j}. \quad (1.3)$$

We can thus see that the metric indeed defines the distance. In the usual Euclidean space, the metric is the identity matrix.

One can then generalise the notion of distance by taking into account the fact that the metric can depend on the position. Thus, similarly to (1.3), we define the interval as

$$ds^2 = dx^\nu dx_\nu = g_{\mu\nu}(x) dx^\mu dx^\nu . \quad (1.4)$$

This equation then tells us that the position does not, by itself, determine the distance anymore, contrarily to Euclidean space! The metric is also partially what makes a spacetime different from another, ‘partially’ because different coordinate systems may represent the same spacetime or portion of it. A change of coordinates indeed changes the expression of the metric. One must also introduce the inverse metric $g^{\mu\nu}$, which is such that

$$g_{\mu\alpha} g^{\mu\beta} = \delta_\alpha^\beta , \quad (1.5)$$

with δ the Kronecker delta.

As already mentioned, in Euclidean space, therefore in Newtonian theory, the metric is represented by the identity matrix. One can show that this space is flat, *i.e.* it has no curvature. The basic flat spacetime is the Minkowski spacetime, which has its metric denoted by

$$\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1) . \quad (1.6)$$

Though the flat spacetime may be easier to manipulate, general relativity predicts that matter has an influence on the spacetime it lies in and curves it. The fact that it can be curved induces some changes in the mathematical tools we use. At each point, we define a tangent space, which contains all the vectors located at that point [17]. Two vectors can be compared in a natural way if they belong to the same tangent space (see [17]). In flat spacetime, all tangent spaces are equivalent, so that it is natural to compare vectors located at different points. This is however not the case in a curved spacetime. In order to compare vectors of different tangent spaces properly, we need to introduce a *connection*, described by the connection coefficients $\Gamma^\alpha_{\mu\nu}$. This connection indicates how one can transport a vector from one tangent space to another. It is used to define the covariant derivative, which takes into account the fact that two adjacent points may not have the same tangent space. This covariant derivative is needed for the derivative operator to remain a tensor, since one can show that the partial derivative is not one. It therefore replaces the partial derivative in a curved spacetime. We define $\partial_\alpha = \partial/\partial x^\alpha$, and we can then express the covariant derivative as

$$\nabla_\alpha V^\mu = \partial_\alpha V^\mu + \Gamma^\mu_{\alpha\lambda} V^\lambda , \quad (1.7)$$

$$\nabla_\alpha W_\mu = \partial_\alpha W_\mu - \Gamma^\lambda_{\alpha\mu} W_\lambda , \quad (1.8)$$

$$\nabla_\alpha q = \partial_\alpha q \quad (1.9)$$

where q is a scalar. This generalises to several indices by adding as many connections as indices and using + or - according to whether the index is at the top or at the bottom.

The connection as it was introduced is not unique, it is simply represented by a set of coefficients that satisfy a given transformation law (under change of coordinates) that

is different from that of tensors and is such that the covariant derivative transforms as a tensor. As a consequence, different connections define different covariant derivatives. We can however get a unique metric-dependent connection if we demand that it have two specific properties. The first is that it must be symmetric, also said torsion-free, and the second is that it must be metric-compatible, *i.e.* the connection is such that the covariant derivative of the metric is null. These properties respectively read

$$\Gamma^\alpha_{\mu\nu} = \Gamma^\alpha_{\nu\mu}, \quad (1.10)$$

$$\nabla_\alpha g_{\mu\nu} = 0, \quad (1.11)$$

It can be shown that imposing these two properties allows one to obtain a unique expression for the connection, which is

$$\Gamma^\alpha_{\mu\nu} = \frac{1}{2} g^{\alpha\beta} (\partial_\mu g_{\nu\beta} + \partial_\nu g_{\mu\beta} - \partial_\beta g_{\mu\nu}). \quad (1.12)$$

The connection defined by this expression is called the *Levi-Civita connection*. In the following, $\Gamma^\alpha_{\mu\nu}$ always refers to the Levi-Civita connection. There are several advantages of using this connection, *e.g.* it appears naturally in the equations of motion, adds symmetries to the Riemann tensor, introduced hereafter, and associates the curvature to the metric itself.

One can show that the curvature of spacetime is related to the Riemann tensor, which is expressed as a function of our connection:

$$R^\alpha_{\beta\mu\nu} = \partial_\mu \Gamma^\alpha_{\nu\beta} - \partial_\nu \Gamma^\alpha_{\mu\beta} + \Gamma^\alpha_{\mu\lambda} \Gamma^\lambda_{\nu\beta} - \Gamma^\alpha_{\nu\lambda} \Gamma^\lambda_{\mu\beta}. \quad (1.13)$$

It is also possible to show that space is flat if the Riemann tensor is null. We can see that the curvature of spacetime depends on the second (and first) derivative(s) of its metric, since we use the Levi-Civita connection. This highlights the fact that the Minkowski metric is flat, since it does not depend on the position, and therefore yields a Riemann tensor with only null elements.

One can further define the Ricci tensor, which is symmetric, and the Ricci scalar respectively as

$$R_{\mu\nu} = R^\alpha_{\mu\alpha\nu}, \quad (1.14)$$

$$R = R^\alpha_{\alpha} = g^{\mu\nu} R_{\mu\nu} (= \text{Tr } R_{\mu\nu}). \quad (1.15)$$

The influence of matter on spacetime is computed from Einstein's equations, which read

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} = 8\pi G T_{\mu\nu}, \quad (1.16)$$

for a zero cosmological constant, with $T_{\mu\nu}$ the energy-momentum tensor, and $G_{\mu\nu}$ the Einstein tensor. In vacuum, $T_{\mu\nu} = 0$. Solving these equations thus allows one to determine the metric of the spacetime.

Finally, there exist three types of trajectories: time-like, null or light-like, and space-like. For a null motion, we have $ds^2 = 0$ and a velocity equal to the speed of light, while time-like motions are characterised by a velocity lower than c and space-like ones larger than c . The path taken by light is described by null curves.

1.2 Cosmological considerations

If one only observes the neighbourhood of the Milky Way, one does not need to take cosmology into account. As the sensitivity of the detectors increases, for both the electromagnetic and gravitational radiations, we can observe further away. One then needs to take into account the effect of the expansion and possible curvature of the Universe. We shall only consider a flat space, but the expansion will be taken into account. I will thus present some cosmological concepts, namely the standard model and the redshift, and discuss the notion of distance. This section is based on [17], [19] and [20].

1.2.1 Cosmological models

The theory of general relativity can be applied to the Universe as a whole, not only to some parts of it. Hence, we can obtain a cosmological model that describes how the Universe evolves with time. The simplest assumption is that, on cosmological scales, the Universe is both homogeneous and isotropic. A famous metric that shows the expansion of the Universe is the Friedmann-Lemaître-Robertson-Walker metric, or FLRW. It is expressed through

$$ds^2 = -dt^2 + a^2(t) (d\vec{r})^2, \quad (1.17)$$

with $a(t)$ the scale factor. It is defined such that $a(t = t_0) = 1$, with t_0 corresponding to the present time (so a is taken relative to the value we would currently measure). This factor $a(t)$ represents the expansion (or contraction) of space, since the measure of spatial distances is stretched by the factor $a(t)$, which can change with time. It was found that the Universe is in expansion. It is possible to write the metric (1.17) in a different form by factoring $a(t)$, which yields

$$ds^2 = a^2(\eta) [-d\eta^2 + (d\vec{r})^2], \quad (1.18)$$

with η the conformal time. It is defined such that

$$d\eta^2 = \frac{dt^2}{a^2(t)} \quad \Rightarrow \quad \eta = \int \frac{dt'}{a(t')}. \quad (1.19)$$

To be more accurate, the metric includes, in fact, a term that takes into account a possible curvature of space. The full metric in spherical coordinates is

$$ds^2 = -dt^2 + a^2(t) \left(\frac{dr^2}{1 - k r^2} + r^2 d\theta^2 + r^2 \sin^2 \theta d\phi^2 \right), \quad (1.20)$$

with k a constant related to the curvature of space. If $k = 0$ then space is flat, corresponding to a Euclidean space, and we retrieve the metric (1.17). If $k < 0$, space is open (negative curvature), and if $k > 0$, space is closed (positive curvature). One can show that these different values lead to very different predictions regarding the age of the Universe, its fate and other characteristics. The current models and observations suggest that space is flat. The standard cosmology also postulates the existence of dark matter and dark energy, the latter emerging from the addition of the cosmological constant Λ in Einstein's equations

(1.16). This dark energy causes the expansion to accelerate, and was also experimentally shown through the observation of high-redshift supernovae [21].

Using the FLRW metric (1.17) and an appropriate energy-momentum tensor (typically for a perfect fluid), we can express the Einstein's equations and obtain the Friedmann-Lemaître equations (here in flat space and without a cosmological constant):

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G \rho}{3}, \quad (1.21)$$

$$\left(\frac{\ddot{a}}{a}\right) = -\frac{4\pi G}{3} (\rho + 3p), \quad (1.22)$$

where the $\dot{}$ represents the time derivative, ρ the density and p the pressure. According to the epoch and model studied, we use different relations between p and ρ . Nowadays, we are in the matter era, *i.e.* matter dominates over radiation (*i.e.* relativistic particles), so that $p \approx 0$.

So, we can solve these equations and find the scale factor. What is more, a is related to the Hubble constant H_0 . The Hubble parameter $H(t)$ is defined as

$$H^2(t) = \left(\frac{\dot{a}}{a}\right)^2. \quad (1.23)$$

The Hubble constant H_0 corresponds to the value $H(t = t_0)$, with t_0 being again the present time. All the variables noted with an index 0 are taken at the present time. This constant marked the experimental discovery that space is expanding, given the Hubble-Lemaître law

$$v(d) = H_0 d, \quad (1.24)$$

where v is the recession velocity of galaxies and d their distance. It was indeed found that most galaxies were receding from us with a velocity described by this law. If H_0 is known, this relation will allow us to measure cosmological distances.

One distinguishes the peculiar velocity of an object, which is its velocity as it moves *through* space, from the recession velocity, which is its velocity as it is carried away *with* space. In other words, the recession velocity is only due to the expansion. Indeed, when space expands, the coordinates of the object, say in the Earth frame of reference, remain exactly the same. However, the way we measure distances changes through $a(t)$, so that the object is getting more and more distant. So, it moves with space because of the expansion, and not through it, since its coordinates remain the same, contrarily to the peculiar motion. This motion can still be interpreted as if the object was going away from us in a non-expanding Universe, leading to the usual Doppler redshift. If we measure this redshift, we can infer the recession velocity.

It may be interesting to note that we can express Eq. (1.21) at t_0 as

$$\rho_0 = \frac{3 H_0^2}{8\pi G}. \quad (1.25)$$

It means that the measurement of H_0 , which we shall discuss later, can lead to the value of the density of matter required to have a flat space (since this is used to write Eq. (1.21)). This value can then be compared to observations to estimate, for example, the current density of dark matter required for space to be flat.

1.2.2 Cosmological redshift

As we have seen, the Universe is expanding. Such an expansion causes the wavelength to increase along with space, so that the wave is redshifted. This is the cosmological redshift and we will see where it comes from.

It is, in fact, a consequence of *cosmological time dilation*. As space is expanding, the distance between a source and an observer increases between the emission of successive photons. As a result, if two photons are emitted at an interval Δt_S at the source, the observer will measure a time interval $\Delta t_O > \Delta t_S$, since the second photon has a larger distance to travel. To be more precise, one has

$$\frac{\Delta t_O}{a(t_O)} = \frac{\Delta t_S}{a(t_S)}. \quad (1.26)$$

Let us consider a wave emitted at a time t_s . Given that $ds^2 = 0$ for light, one can write, from the FLRW metric,

$$\frac{dt}{a(t)} = dr. \quad (1.27)$$

In the reference frame of the source, the time required to travel a *Euclidean distance* Δr is thus $\Delta t_s/a(t_s)$, assuming $a(t)$ is constant in that time interval. In the reference frame of the observer, the time needed to travel the same Euclidean distance is $\Delta t_0/a_0$. Since the Euclidean distances are the same, one can equal the two expressions, which leads to the result (1.26) (see [20] for a more rigorous proof).

As consequence, if the time between two crests of the wave at emission is Δt_s , it corresponds to a wavelength $\lambda_s = \Delta t_s$ (remember that we use $c = 1$). The observed time between two crests is then $\Delta t_0 = \Delta t_s/a_s$, with $a_s \equiv a(t_s)$, corresponding to an observed wavelength $\lambda_0 = \Delta t_0$. As a consequence, one has

$$\lambda_s = a_s \lambda_0, \quad (1.28)$$

where λ_s and λ_0 are then the emitted and observed wavelengths, respectively. Hence, $\lambda_0 > \lambda_s$ (since a increases with time and $a_0 = 1$), which means that the wavelength increases because of the expansion. This is called *cosmological redshift*.

The redshift is defined as

$$z = \frac{\lambda_0 - \lambda_s}{\lambda_s}, \quad (1.29)$$

We can then inject the result (1.28) into the definition of the redshift to get

$$z = \frac{\lambda_0 - \lambda_s}{\lambda_s} = \frac{a_0 - a_s}{a_s}. \quad (1.30)$$

We know that $a_0 = 1$, so that the relation becomes

$$z = \frac{1 - a}{a} \quad (1.31)$$

$$a = \frac{1}{1 + z} \quad (1.32)$$

1.2.3 Distances in cosmology

Because of the expansion, the notion of distances may become somewhat complex. I will introduce a few distances that are used in cosmology.

The first one is the proper distance, which corresponds to the distance between two objects at a given fixed time. This seems obvious and corresponds to the usual distance we use, but it cannot be used here in practice. Indeed, we do not have any ruler to measure instantaneously the distance to an object at cosmological scales. It is not the distance we measure using the light we receive, since the object will have moved during the photon travel time. In a flat Universe and at present time, it corresponds to the *comoving distance*, *i.e.* the Euclidean distance. We can write it at any time with

$$d_p(t) = a(t) r_0 = \frac{r_0}{1+z}, \quad (1.33)$$

where r_0 is the comoving distance. The equation thus represents the proper distance of an object at a redshift z at the time it emitted the light we receive. Note that the comoving distance cannot be measured either.

Then, there is the luminosity distance, which is the distance we deduce from the known luminosity of an object and the received flux assuming it is only subjected to geometric dilution. In other words, we have

$$d_L^2 \equiv \frac{L}{4\pi F}. \quad (1.34)$$

To relate it to the proper distance, we need to take into account the cosmological redshift that adds to the geometric dilution. First, the wavelength increases by a factor $(1+z)$ due to the expansion of space, so that the energy and hence the flux is reduced by the same amount. Then, the cosmological time dilation causes another reduction of the flux by a factor $(1+z)$. As a consequence, we can write

$$F = \frac{L}{4\pi r_0^2 (1+z)^2} \quad (1.35)$$

$$\Rightarrow d_L = r_0 (1+z) \quad (1.36)$$

Finally, we can introduce the angular diameter distance, which is the one that we will use in the gravitational lensing framework. This distance relates the diameter D of an object to its angular size θ , *i.e.*

$$d_A \equiv \frac{D}{\theta}. \quad (1.37)$$

We can see that it is equivalent to a Euclidean distance for which we would write

$$\tan \theta = \frac{D}{d} \Rightarrow \theta \approx \frac{D}{d}, \quad (1.38)$$

since the angles are very small, given the distance scale in cosmology. The angular diameter distance can be related to the comoving and luminosity distances. It is possible to show that

$$d_A = \frac{d_L}{(1+z)^2} = \frac{r_0}{1+z}. \quad (1.39)$$

1.3 Summary of conventions and important relations

A summary of the important relations and conventions used throughout the whole thesis is presented in Tab. 1.1.

Curvature of space	null, flat space
Units	$c = 1$
Minkowski metric	$\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$
Levi-Civita connection	$\Gamma^{\alpha}_{\mu\nu} = \frac{1}{2}g^{\alpha\beta}(\partial_{\mu}g_{\nu\beta} + \partial_{\nu}g_{\mu\beta} - \partial_{\beta}g_{\mu\nu})$
Riemann tensor	$R^{\alpha}_{\beta\mu\nu} = \partial_{\mu}\Gamma^{\alpha}_{\nu\beta} - \partial_{\nu}\Gamma^{\alpha}_{\mu\beta} + \Gamma^{\alpha}_{\mu\lambda}\Gamma^{\lambda}_{\nu\beta} - \Gamma^{\alpha}_{\nu\lambda}\Gamma^{\lambda}_{\mu\beta}$
Ricci tensor, Ricci scalar	$R_{\mu\nu} = R^{\alpha}_{\mu\alpha\nu}, R = R^{\alpha}_{\alpha}$
Einstein's equation	$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}Rg_{\mu\nu} = 8\pi GT_{\mu\nu}$
FLRW metric (flat space)	$ds^2 = -dt^2 + a^2(t)(d\vec{r})^2$
Scale factor - redshift	$a = \frac{1}{1+z}$
Angular diameter distance	$d_A = \frac{r_0}{1+z}$

Table 1.1: Summary of the conventions and important relations

Chapter 2

Electromagnetic gravitational lensing

In this chapter, we will introduce the concept of gravitational lensing and how it is described in the case of a luminous signal. As we shall see later, this description will be useful in the case of gravitational waves. This whole chapter is largely based on [22], [23] and [24]. I invite the interested reader to have a look at these references for more details.

2.1 Gravitational lensing effect

As we have seen in the previous section, matter deforms spacetime and objects that are massive enough may produce a significant curvature. This has a very important consequence: gravity also influences massless bodies, while this is not the case in the Newtonian theory. Indeed, we must not think of gravity as a ‘force’ acting on massive bodies. Since the very space on which any particle moves is curved, its trajectory will inevitably be influenced by this curvature, whether it has a mass or not.

As a result, we find that massive bodies can deflect not only matter but also light. This phenomenon is at the core of the gravitational lensing effect. To be more precise, we speak of *gravitational lensing* when we observe one or several deformed and/or magnified image(s) of the same object. The apparent position of the source is thus different from its true location. This is possible because of this deflection of light by massive bodies. The deflecting body is called the *lens*, or the deflector. The properties of the images strongly depend on the lens, so we can retrieve information on the latter by observing this effect. Its characteristics and interests will be developed later.

We distinguish two different types of gravitational lensing: strong and weak. *Strong lensing* occurs when there are several images. We talk about macrolensing when the multiple images are resolved, so that we can tell them apart, while we talk about microlensing when they are unresolved. In the former case, the lens is typically a galaxy or galaxy cluster, while it is typically a star in the second case. As for *weak lensing*, the lens does not allow for multiple images, but still distorts the image. This type of lensing is more complicated to detect since we cannot simply tell whether what we see is the true shape or a distorted view of the source.

2.2 Lens equation

We will now describe the phenomenon mathematically in flat space. This section is largely based on [22]. We make the thin lens approximation, *i.e.* we assume that the dimensions of the lens are small compared to the considered distances. As a consequence, we can assume that the lens lies on a plane called the lens plane, which is the plane on which the images ‘form’. Similarly, the source lies on the source plane. The phenomenon is thus described using 2D vectors lying in these two planes. The situation we consider is represented in Fig. 2.1.




Figure 2.1: Sketches of the geometry of the problem (a) in 3D (b) in the plane containing the light ray. Dimensions are not to scale, angles are exaggerated compared to distances. Note that for any quantity \vec{x} , one defines $x \equiv |\vec{x}|$.

The observer is denoted by O , the lens by L and the source by S . In Fig. 2.1, we consider that the plane containing the ray also contains O , L and S ¹. The reference axis from which one measures angles is chosen to be the one crossing O and L . The source is thus at an angular position $\vec{\beta}$, while the apparent source (in the source plane) S' is at an angle $\vec{\theta}$. If there is no lens, $\vec{\beta} = \vec{\theta}$, since we see the source where it really is as there is no deflection.

¹In a more general case, there may be two different planes, one containing the ray between L and S , and a different one between O and L .

Note that we use 2D vectors to describe these angles. These are such that the norm of the vector is the value of the angle in the plane of the ray, *i.e.* $|\vec{\theta}| = \theta$, and its direction is parallel to the intersection of the plane of the ray and a plane perpendicular to it (*e.g.* the lens plane). For example, as we shall see, $\vec{\theta}$ is parallel to $\vec{\xi}$ and $\vec{\beta}$ to $\vec{\eta}$.

We also define the deflection angle $\hat{\alpha}$ as being the angle between the trajectory taken by the deflected light ray and the trajectory it would have if it were not deflected. In this work, we will only consider single-plane lensing, which means that we assume that there is only one lens, and thus one lens plane. It can be generalised with several lenses, hence several light deflections. We also consider small angles, given the distance scales.

The *angular diameter distance* from the observer to the lens plane is D_L , to the source D_S and the one between the lens and the source is D_{LS} . In general (at cosmological scales), $D_S \neq D_L + D_{LS}$ given the type of distance we use. Note that we approximate the trajectory by neglecting the actual deviation near the lens plane, given the small-angle approximation. In other words, we consider the trajectory as made of two asymptotes, *i.e.* straight lines, while the actual trajectory is curved.

In the plane of the lens, we define the impact parameter of the deflected light ray $\vec{\xi}$, the norm of which represents the minimal distance between the lens and the light ray. We can express it using the definition of the angular diameter distance as $\vec{\xi} = D_L \vec{\theta}$. Intuitively, this comes from the projections of the ray and $\vec{\xi}$ onto planes defined by the reference axis and another one perpendicular to OL . This is depicted in Fig. 2.2. Using angular diameter distances, we can see that $\xi_1 = D_L \theta_1$ and $\xi_2 = D_L \theta_2$. This is summarised with vectors as $\vec{\xi} = D_L \vec{\theta}$, with $\vec{\xi} = (\xi_1, \xi_2)$ and $\vec{\theta} = (\theta_1, \theta_2)$ (in the axes \vec{e}_1 and \vec{e}_2 , see Fig. 2.2). The vectors to describe angles are thus useful. They allow us to consider the very position at which the ray crosses the lens plane and not only its distance to L (*i.e.* with only $\xi = D_L \theta$), which is important if the problem is not axially symmetric.




Figure 2.2: Projections of the ray and $\vec{\xi}$ onto two different reference planes. (a) Definition of the axes \vec{e}_1 and \vec{e}_2 on the lens plane. (b) Projection onto the plane defined by OL and \vec{e}_1 (c) Projection onto the plane defined by OL and \vec{e}_2 .

Also, we represent the misalignment of the source with respect to the lens and the observer by $\vec{\eta}$, the norm of which is the distance to the reference axis. Once again we can use the angles to express it as $\vec{\eta} = D_S \vec{\beta}$. The angle $\vec{\theta}$ thus indicates the angular position of the source such that we see it because of the light deflection, so its apparent position, while $\vec{\beta}$ denotes the true position. Intuitively, we can say that the deflection angle depends on $\vec{\theta}$, *i.e.* $\hat{\alpha} = \hat{\alpha}(\vec{\theta})$. Indeed, the effects of the mass of the lens get lower as we go away from the

lens so that the deflection will be weaker as we go to larger impact parameters. The relation between $\hat{\alpha}(\vec{\theta})$ and $\vec{\theta}$ depends on the model that is used to describe the lens.

We can relate the three angles by a simple relation, provided they are small. We have that the vector $\vec{\eta}'$ that connects the apparent source and the reference axis is

$$\vec{\eta}' = D_S \vec{\theta}. \quad (2.1)$$

This vector can be divided into two contributions, namely

$$\vec{\eta}' = \vec{\eta} + \vec{\eta}_{SS'}, \quad (2.2)$$

where $\vec{\eta}_{SS'}$ is the vector between S and S' on the plane of the source. We can use once again the definition of the angular distance, with $\vec{\eta}$ related to $\vec{\beta}$ and $\vec{\eta}_{SS'}$ to $\hat{\alpha}$, to write:

$$D_S \vec{\theta} = D_S \vec{\beta} + D_{LS} \hat{\alpha}. \quad (2.3)$$

We can express it in a different manner by dividing by D_S and introducing the reduced deflected angle

$$\vec{\alpha}(\vec{\theta}) \equiv \frac{D_{LS}}{D_S} \hat{\alpha}(\vec{\theta}). \quad (2.4)$$

Thus, Eq. (2.3) becomes

$$\vec{\beta} = \vec{\theta} - \vec{\alpha}(\vec{\theta}), \quad (2.5)$$

which is called the *lens equation*. Note that the deflection angle $\vec{\alpha}$ does not depend on $\vec{\beta}$, since the initial direction of the light ray does not matter, it will be deflected by the same amount at the same impact parameter. What changes with $\vec{\beta}$ is the deflection angle required to reach the observer, and so the angle $\vec{\theta}$. Note also that α cannot be large if both β and θ are small

Therefore, this equation allows us to determine where the object will be observed ($\vec{\theta}$), according to its true position relative to the lens ($\vec{\beta}$), or determine its true position from the images that are observed. We can see that the equation will strongly depend on the expression of $\vec{\alpha}(\vec{\theta})$, so that a critical point is to model the deflection angle correctly.

The lens equation can be modified further by introducing dimensionless quantities and expressing the angles using angular distances. Thus, we introduce

$$\vec{x} \equiv \frac{\vec{\xi}}{\xi_0} = \frac{D_L \vec{\theta}}{\xi_0}, \quad (2.6)$$

$$\vec{y} \equiv \frac{\vec{\eta}}{\eta_0} = \frac{D_S \vec{\beta}}{\eta_0}, \quad (2.7)$$

$$\eta_0 = \frac{D_S}{D_L} \xi_0; \quad (2.8)$$

where ξ_0 and η_0 are length scales in the lens and source planes, respectively. The expression of η_0 is given to make both length scales correspond to the same characteristic angle. Note that if $\xi_0 = D_L$, then $\vec{x} = \vec{\theta}$ and $\vec{y} = \vec{\beta}$.

If we use the last three relations in the lens equation, it becomes

$$\begin{aligned} \frac{D_S}{D_L} \xi_0 \frac{1}{D_S} \vec{y} &= \xi_0 \frac{1}{D_L} \vec{x} - \vec{\alpha}(\xi_0 \vec{x}) \\ \Rightarrow \vec{y} &= \vec{x} - \vec{\alpha}(\vec{x}), \end{aligned} \quad (2.9)$$

where we have introduced a scaled deflection angle

$$\vec{\alpha}(\vec{x}) = \frac{D_L}{\xi_0} \frac{D_{LS}}{D_S} \hat{\alpha}(\xi_0 \vec{x}). \quad (2.10)$$

Finally, it is important to keep in mind that all the assumptions that were made, namely a thin lens, small angles and the approximation of the path by straight lines, will also hold in the rest of Section 2, unless explicitly specified.

2.3 Time delays

Since the deflected rays do not follow the same path as unlensed ones, they need more time to reach an observer. In other words, the lens introduces a delay in the arrival time. I shall present the main results and you may find a more detailed development in Appendix A.1 or cited references.

One can compute the delay in a Minkowski background metric perturbed by a weak lens. It is expressed through [17, 22]

$$ds^2 = -(1 + 2U) dt^2 + (1 - 2U) (d\vec{r})^2, \quad (2.11)$$

where \vec{r} is the spatial position vector and $U \equiv U(\vec{r}) = -GM/r$ is the gravitational potential. This metric assumes that $U \ll 1$ (or $U \ll c^2$ in SI units), which is called the weak field limit. Since we do not consider the expansion yet, the distances are Euclidean ones and I will refer to them as d instead of D to avoid the confusion with angular diameter distances. As light follows null geodesics, we can write, from the metric,

$$ds^2 = -(1 + 2U) dt^2 + (1 - 2U) (d\vec{r})^2 = 0 \quad (2.12)$$

$$\Rightarrow dt^2 = \frac{1 - 2U}{1 + 2U} (d\vec{r})^2 \quad (2.13)$$

$$\Rightarrow dt = \sqrt{\frac{1 - 2U}{1 + 2U}} |d\vec{r}| \quad (2.14)$$

We introduce $dl = |d\vec{r}|$, with dl the arc length in Euclidean space (since in that case in Cartesian coordinates $dl^2 = dx^2 + dy^2 + dz^2 = (d\vec{r})^2$). We can express the potential term to first order, which leads to

$$\sqrt{\frac{1 - 2U}{1 + 2U}} = 1 - 2U + O(U^2). \quad (2.15)$$

If we further consider the source to emit at a time $t = 0$, we get:

$$t \simeq \int (1 - 2U) dl = l - 2 \int U dl , \quad (2.16)$$

with l the Euclidean length of the path, and the potential is integrated over that path. This equation thus expresses the time it takes for a light ray to reach the observer. We can see that it has two contributions: a geometrical one, represented by the first term, and a gravitational one, caused by the gravitational potential U . This term is called the Shapiro time delay. It indeed increases t , since U is negative.

One can show (see Appendix A.1) that the time delay Δt , which is the difference between the arrival times of the lensed ray and the unlensed one, is

$$\Delta t(\vec{x}) = \frac{d_S \xi_0^2}{d_L d_{LS}} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}) \right] , \quad (2.17)$$

where we define the dimensionless lensing potential

$$\psi(\vec{x}) = 2 \frac{d_{LS} d_L}{d_S \xi_0^2} \int U(\vec{x}, Z) dZ . \quad (2.18)$$

We have used $\vec{r} = \vec{x} + Z\vec{e}_Z$ with \vec{e}_Z the vector perpendicular to the lens plane and such that $U(\vec{r}) = U(\vec{x}, Z)$. Note that the integral is not done over dl anymore, contrarily to Eq. (2.16). As the deflection angle is considered to be small, one can approximate the integral over the true path by the integral over the path of an unlensed ray, with the third dimension Z taken along this ray.

The lensing, or deflection, potential $\psi(\vec{x})$ is a projection of the potential onto the lens plane, as U is integrated over the third dimension (outside the lens plane). The source of this potential is the projected surface density

$$\Sigma(\vec{x}) = \int \rho(\vec{x}, Z) dZ . \quad (2.19)$$

We can also consider $\xi_0 = d_L$, which yields

$$\psi(\vec{x}) = 2 \frac{d_{LS}}{d_S d_L} \int U(\vec{x}, Z) dZ , \quad (2.20)$$

$$\Delta t(\vec{x}) = \frac{d_S d_L}{d_{LS}} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}) \right] . \quad (2.21)$$

We have considered a spatial Euclidean space, but in practice, things are a bit more complicated. We need to consider, in the most general case, a cosmological model that takes into account all the large-scale effects such as the expansion of space. As a result, the background metric can be taken to be the FLRW metric, as introduced in Eq. (1.17). However, the metric is perturbed by the presence of a lens the same way it was for a Minkowski background metric. We thus have [23, 24]

$$ds^2 = -(1 + 2U) dt^2 + a^2(t) (1 - 2U) (d\vec{r})^2 . \quad (2.22)$$

It can be shown (see Appendix A.1) that, in this metric, we have

$$\psi(\vec{x}) = 2 \frac{D_{LS}}{D_S D_L} \int U(\vec{x}, Z) dZ, \quad (2.23)$$

$$\Delta t(\vec{x}) = (1 + z_L) \frac{D_S D_L}{D_{LS}} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}) \right]. \quad (2.24)$$

We see that the definition of $\psi(\vec{x})$ has not changed, except for the type of distance that is used, while the total time delay is simply multiplied by a factor $(1 + z_L)$. We can also introduce *Fermat's potential* $\Psi(\vec{x})$, defined as

$$\Psi(\vec{x}) = \frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}), \quad (2.25)$$

i.e. the time delay without the multiplicative constants.

We can also compute the relative time delay between two images, which is simply the difference between the two respective time delays of these images, since these are computed relative to the same unlensed ray. For a more rigorous and complete development, I recommend having a look at [23] and [24].

2.4 Deflection angle

As already stated, the expression of the deflection angle is very important in the solution of the lens equation. We will now show that it is possible to retrieve it from the time delay. The computation of the latter also allows one to retrieve the lens equation through Fermat's principle. This principle states that a null curve going from a source S to a time-like observer O is a light ray if this curve has a stationary value of the arrival time, as measured by the observer [23]. In other words, if we consider all the null curves going from S to O , each with its own arrival time t , the light ray is the one for which t is an extremum of all the t considered [18].

We can naturally substitute the arrival time (\vec{x}) for the time delay, since $\Delta t(\vec{x}) = t(\vec{x}) - t_u$, where t_u is the arrival time of the unlensed ray and is constant (or such that $\delta t_u = 0$ by Fermat's principle). It will therefore vanish upon differentiation. As we have computed the time delay at the observer, we now demand that it be an extremal value of all the possible time delays. The condition reads

$$\delta [\Delta t(\vec{x})] = 0, \quad (2.26)$$

or equivalently

$$\vec{\nabla}_{\vec{x}} [\Delta t(\vec{x})] = 0. \quad (2.27)$$

We can use its expression from Eq. (2.24) to get

$$\vec{\nabla}_{\vec{x}} [\Delta t(\vec{x})] = (\vec{x} - \vec{y}) - \vec{\nabla}_{\vec{x}} \psi(\vec{x}) = 0. \quad (2.28)$$

If we use the lens equation (2.9), we finally arrive at

$$\vec{\alpha}(\vec{x}) = \vec{\nabla}_{\vec{x}} \psi(\vec{x}). \quad (2.29)$$

We can then see that the deflection angle can be expressed in terms of the lensing potential and that the lens equation is in agreement with Fermat's principle². What is interesting, then, is that the number of extrema (and saddle points) of the time delay gives the number of images produced by the lensing effect.

If we further develop the expression of the scaled deflection angle using Eq. (2.23), we get

$$\begin{aligned}\vec{\alpha}(\vec{x}) &= \vec{\nabla}_{\vec{x}} \psi(\vec{x}) = \vec{\nabla}_{\vec{x}} \left[2 \frac{D_{LS}}{D_S D_L} \int U(\vec{x}, Z) dZ \right] \\ &= 2 \frac{D_{LS}}{D_S D_L} \int \vec{\nabla}_{\vec{x}} U(\vec{x}, Z) dZ\end{aligned}\quad (2.30)$$

If one uses the refraction index and Fermat's principle on the spatial path of light rays ([22], [23], [24]), the deflection angle that is found is

$$\vec{\alpha}(\vec{x}) = 2 \frac{D_{LS}}{D_S} \int \vec{\nabla}_{\perp} U(\vec{x}, Z) dZ, \quad (2.31)$$

where $\vec{\nabla}_{\perp}$ refers to the gradient projected onto the plane perpendicular to the light ray. In fact, given the small angles (in particular here $\hat{\alpha}$) and the thin lens assumption, we can approximate the plane perpendicular to the light ray by the lens plane, *i.e.* approximate the operator $\vec{\nabla}_{\perp}$ by the gradient in the lens plane. As a consequence, we can rewrite Eq. (2.30) as

$$\vec{\alpha}(\vec{x}) = 2 \frac{D_{LS}}{D_S D_L} \int \vec{\nabla}_{\vec{x}} U(\vec{x}, Z) dZ = 2 \frac{D_{LS}}{D_S} \int \vec{\nabla}_{\perp} U(\vec{x}, Z) dZ, \quad (2.32)$$

and we retrieve the expression (2.31). The cancellation of D_L comes from the fact that $\vec{\nabla}_{\perp}$ is not dimensionless, contrarily to $\vec{\nabla}_{\vec{x}}$ ($= \xi_0 \vec{\nabla}_{\vec{\xi}}$, with $\xi_0 = D_L$).

We can express the deflection for a point-mass lens (see Appendix A.2), which is

$$\hat{\alpha}(\vec{\xi}) = 4GM \frac{\vec{\xi}}{|\vec{\xi}|^2}, \quad (2.33)$$

$$|\hat{\alpha}(\vec{\xi})| = \frac{4GM}{\xi}, \quad (2.34)$$

with $\xi = |\vec{\xi}|$ the impact parameter. The most general lens can be considered as made of a lot of point masses. The contribution of each point mass adds to the other, as the deflection angle is linear in M . We can thus write

$$\hat{\alpha}(\vec{\xi}) = \sum_i \hat{\alpha}(\vec{\xi} - \vec{\xi}_i) = 4G \sum_i \frac{M_i (\vec{\xi} - \vec{\xi}_i)}{|\vec{\xi} - \vec{\xi}_i|^2}, \quad (2.35)$$

²The development was found to be made the other way around, *i.e.* the relation (2.29) is derived using the refraction index and Fermat's principle. If we change $\vec{\nabla}_{\vec{x}} \psi(\vec{x})$ into $\vec{\alpha}(\vec{x})$ in Eq. (2.28), we retrieve the lens equation, so we see that Fermat's principle directly leads to the lens equation.

where $\vec{\xi}$ is still the impact parameter vector and $\vec{\xi}_i$ denotes the position of the point mass i , such that the impact parameter of a point mass i is $(\vec{\xi} - \vec{\xi}_i)$. In the continuous limit, one would need to use the density instead of the mass and replace the sum by an integral on the lens plane. However, given the thin lens approximation, the mass distribution of the lens is only described by its surface density $\Sigma(\vec{\xi})$ given in equation (2.19). We thus have

$$\Sigma(\vec{\xi}) = \int \rho(\vec{\xi}, Z) dZ, \quad (2.36)$$

$$\hat{\alpha}(\vec{\xi}) = 4G \int \frac{\Sigma(\vec{\xi}') (\vec{\xi} - \vec{\xi}')}{|\vec{\xi} - \vec{\xi}'|^2} d\vec{\xi}'. \quad (2.37)$$

We can see that the integration is performed on the lens plane and not on the third dimension anymore.

We can use a similar development for the lensing potential to express it as an integral over the lens plane. It then reads

$$\psi(\vec{\xi}) = 4G \int \Sigma(\vec{\xi}') \ln |\vec{\xi} - \vec{\xi}'| d\vec{\xi}'. \quad (2.38)$$

2.5 Magnification and distortion

As already mentioned, gravitational lensing can magnify, reduce and/or distort images because the deflection depends on the distance to the lens, *i.e.* the impact parameter. As a consequence, two neighbouring rays will not be deflected in the same way. The developments that will be presented are valid under the assumption that the angular size of the object is small compared to the angular scale over which the properties of the lens change.

It appears that the surface brightness is conserved. Therefore, the flux, which is the product of the surface brightness with the solid angle, is magnified or demagnified according to the change in solid angle. The magnification μ of the image of a small extended source can then be defined as [23, 24]

$$\mu = \frac{d\Omega_L}{d\Omega_S} = \frac{D_S^2 dA_L}{D_L^2 dA_S}, \quad (2.39)$$

with $d\Omega_L$ the infinitesimal solid angle subtended by the image, $d\Omega_S$ the one subtended by the image of the source if there were no lensing effect, and dA_L and dA_S the corresponding areas. It can be shown that one can express the magnification with the Jacobian matrix $A(\vec{x})$, defined as

$$A(\vec{x}) = \frac{\partial \vec{y}}{\partial \vec{x}} \quad \text{or} \quad A_{ij} = \frac{\partial y_i}{\partial x_j}. \quad (2.40)$$

This matrix represents the mapping from \vec{x} to \vec{y} described by the lens equation. In fact, the determinant of A corresponds to the ratio of the infinitesimal area on the source plane over

its corresponding area on the lens plane through the mapping $\vec{x} \rightarrow \vec{y}$ [23]. Given Eq. (2.39), the magnification can thus be expressed as

$$\mu = \frac{1}{\det A}. \quad (2.41)$$

The lens equation (2.9) provides the mapping from \vec{x} to \vec{y} described by A , so that we can write

$$A_{ij} = \delta_{ij} - \frac{\partial^2 \psi(\vec{x})}{\partial x_i \partial x_j}, \quad (2.42)$$

given that $\vec{\alpha} = \vec{\nabla} \psi$. To simplify the notation, I shall refer to $\frac{\partial^2 \psi(\vec{x})}{\partial x_i \partial x_j}$ as ψ_{ij} .

We can introduce a couple of interesting quantities. The first one is the convergence κ , defined in its dimensionless version as

$$\kappa(\vec{x}) = \frac{\Sigma(\vec{x})}{\Sigma_{cr}}, \quad (2.43)$$

with Σ_{cr} the critical surface density

$$\Sigma_{cr} = \frac{1}{4\pi G} \frac{D_S}{D_{LS} D_L}. \quad (2.44)$$

This density is critical in the sense that it is a sufficient condition (but not necessary) to have multiple images (see [24]).

Another interesting thing to notice is that the determinant of A can be null, which means that the magnification can be infinite. The sets of points at a location \vec{x} such that $\det A = 0$ are called *critical curves*. The corresponding curves on the source plane are called *caustics*. Thus, these are obtained through the mapping from \vec{x} to \vec{y} (again with the lens equation). In other words, the caustics represent the positions of the source where the magnification of its image through gravitational lensing is formally infinite. In practice, the magnification will never be infinite, since this would require the source to be a point (given Eq. (2.39)), which is never perfectly the case.

Finally, the total magnification of an object is defined as the sum of the absolute value of the magnification of each image.

2.6 Types and number of images

As we have seen, the lens equation is equivalent to Fermat's principle $\vec{\nabla} [\Delta t(\vec{x})] = 0$, which means that images will form on extrema and saddle points of the time delay surface $\Delta t(\vec{x})$. To know whether the image corresponds to a minimum, maximum or saddle point, we need to compute the Hessian matrix of the time delay, namely

$$T_{ij} = \frac{\partial^2 \Delta t(\vec{x})}{\partial x_i \partial x_j}, \quad (2.45)$$

or, if we use the expression of the gradient shown in Eq. (2.28),

$$\begin{aligned} T &= \vec{\nabla}_{\vec{x}} \left(\vec{\nabla}_{\vec{x}} \Delta t \right) = (1 + z_L) \frac{D_S D_L}{D_{LS}} \vec{\nabla}_{\vec{x}} \left[(\vec{x} - \vec{y}) - \vec{\nabla}_{\vec{x}} \psi(\vec{x}) \right] \\ &= (1 + z_L) \frac{D_S D_L}{D_{LS}} \left(\mathbb{1}_2 - \vec{\nabla}_{\vec{x}} \vec{\nabla}_{\vec{x}} \psi(\vec{x}) \right), \end{aligned} \quad (2.46)$$

where $\mathbb{1}_2$ is the two-dimensional identity matrix. Note also that there is no scalar product between the two $\vec{\nabla}$ operators. This can be expressed as

$$\left(\vec{\nabla}_{\vec{x}} \vec{\nabla}_{\vec{x}} \psi(\vec{x}) \right)_{ij} = \frac{\partial^2 \psi(\vec{x})}{\partial x_i \partial x_j} = \psi_{ij}. \quad (2.47)$$

As a result, we can see that

$$T_{ij} = (1 + z_L) \frac{D_S D_L}{D_{LS}} (\delta_{ij} - \psi_{ij}) \propto A, \quad (2.48)$$

with A the Jacobian matrix defined in Eq. (2.42). As the multiplying constants are strictly positive, the ‘sign’ of the Hessian matrix, determining the type of extremum, is the same as that of the Jacobian matrix A .

From these results, we can define three types of images [22]:

- **Type I images:** these appear at a minimum of the time delay surface. Therefore, this requires that the Hessian matrix (and hence A) be positive-definite, *i.e.* that all its eigenvalues be positive (since it is a symmetric matrix). For that, we need $\det A > 0$ and $\text{Tr } A > 0$. Given Eq. (2.41) and since $\det A > 0$, the magnification is also positive (the image is not reversed).
- **Type II images:** they appear at a saddle point of the time delay surface. In that case, A must be indefinite and have eigenvalues of opposite signs, *i.e.* $\det A < 0$. The magnification is therefore negative, and the image is reversed.
- **Type III images:** they appear at a maximum of the time delay surface, which means that A is negative-definite. This requires $\det A > 0$ and $\text{Tr } A < 0$, so that the magnification is positive.

To be more precise, I refer to a ‘reversed image’ by stating that its parity is negative, while the parity of an image that is not ‘reversed’ is positive. A negative magnification therefore means that the parity is flipped.

As we already mentioned, the time delay surface, through the number of its extrema and saddle points, completely determines the number of images produced by the lensing effect. For example, if there is no lens, the lensing potential is null and the time delay is simply a parabola centred on $\vec{x} = \vec{y}$. This means that there is a minimum at the position of the source, and the image that we see obviously keeps its form and parity.

As another example, in the case of a perfect alignment between O , L and S , and with an axially symmetric lens, the image that is observed is a ring. If we consider a point-mass lens, we have $\vec{y} = 0$ and $\psi(\vec{x}) \propto \ln|\vec{x}|$ (since $\Sigma(\vec{x}) = \delta(\vec{x})$), so that we have from Eq. (2.24)

$$\Delta t(\vec{x}) \propto \frac{1}{2}(\vec{x})^2 - \psi(\vec{x}), \quad (2.49)$$

and thus

$$\Delta t(\vec{x}) = \Delta t(-\vec{x}) = \Delta t(|\vec{x}|). \quad (2.50)$$

As a consequence, there is an axial symmetry of the time delay around the lens, it is the same in each direction at a given impact parameter, so any extremum (in fact minimum here) will appear in a ring called Einstein ring. When the misalignment increases, the symmetry of the time delay surface is broken and the surface deforms. At some point, it changes so much that a maximum and a saddle point merge, making two images disappear [22].

The number and types of images are described in detail in [24], and I will here summarise the results. Under the assumption that the surface density is smooth and decreases faster than $|\vec{x}|^2$ for large $|\vec{x}|$, one can show that the number of images is always finite and that there is always at least one image of type I (a single image of type I being the limit of no lensing). Moreover, if we define n_i the number of images of type i , we have

$$n_I + n_{III} = 1 + n_{II} \quad (2.51)$$

$$\Rightarrow n = n_I + n_{II} + n_{III} = 1 + 2n_{II}. \quad (2.52)$$

This tells us that the total number of images n is always odd³ and that, if there are multiple images, there must always be at least an image of type III or one of type II. In fact, for an isolated lens, there can be multiple images if and only if there is an image of type II. Indeed, it requires that $\det A < 0$. If $\det A > 0$ for all \vec{x} , then we can invert the mapping described by A (since it is never null) so that there is no \vec{x} related to the same \vec{y} , *i.e.* there is only one image. On the other hand, if $\det A < 0$ at \vec{x}_0 , then the image is of type II and, according to (2.52), there must be two other images.

Furthermore, one can establish that a convergence κ greater than 1, *i.e.* a surface density $\Sigma(\vec{x})$ greater than the critical surface density Σ_{cr} (2.44), is a sufficient condition (but not necessary) for having multiple images.

Finally, it is also possible to show that the number of images changes by two units when a source crosses a caustic. In the lens plane, two couples of images each merge into one image on the critical curves, leading to a loss of two images. Similarly, one image can be split into two images.

2.7 Point-mass lens

As already mentioned, it is important to model correctly the deflection angle. To this end, one needs a correct representation of the lens and its mass distribution. There thus

³One often only sees an even number of images, but there can be a highly demagnified image close to the lens that can hardly be observed.

exist several lens models and I will present the most simple one: the point-mass lens. This section is largely based on [22].

The deflection angle has already been defined for a point mass in Eq. (2.33), so that we can write the lens equation using β and θ . We can drop the vector notation given the axial symmetry of the problem. It does indeed not matter whether the source is above or below the reference axis, only its distance to it is important. We can thus write

$$\beta = \theta - 4GM \frac{D_{LS}}{D_S} \frac{1}{D_L \theta} = \theta - \frac{\theta_E^2}{\theta}, \quad (2.53)$$

where we introduced the *Einstein radius*

$$\theta_E = \sqrt{4GM \frac{D_{LS}}{D_S D_L}} \quad (2.54)$$

corresponding to the *linear Einstein radius*

$$R_E = D_L \theta_E. \quad (2.55)$$

We can then obtain a quadratic equation

$$\theta^2 - \beta \theta - \theta_E^2 = 0, \quad (2.56)$$

which can be rewritten by introducing $x = \theta/\theta_E$ and $y = \beta/\theta_E$

$$x^2 - yx - 1 = 0. \quad (2.57)$$

The solution is quite straightforward and reads

$$x_{\pm} = \frac{1}{2} \left(y \pm \sqrt{y^2 + 4} \right). \quad (2.58)$$

We can see that if $y = 0$, we get

$$x_{\pm} = \pm 1, \quad (2.59)$$

which means that both images are on a ring of radius θ_E , as already mentioned earlier. One can also note that $x_- < 0$ and $x_+ > 0$, so that they are always on opposite sides. Finally, we can see that for $y \gg 1$, $x_+ \rightarrow y$ and $x_- \rightarrow 0$, *i.e.* there is no lensing.

It can be shown (see Appendix A.3) that the magnification is

$$\mu_{\pm} = \frac{1}{2} \pm \frac{y^2 + 2}{2y \sqrt{y^2 + 4}}. \quad (2.60)$$

As a result, the $+$ image is always magnified, since $\mu_+ \geq 1$ for all y , while the $-$ image can be magnified or demagnified, as $\mu_- \leq 0$ for all y , so $|\mu_-|$ can be smaller or greater than 1. The total magnification of a point source can then be calculated with

$$\mu = |\mu_+| + |\mu_-| = \frac{y^2 + 2}{y \sqrt{y^2 + 4}}. \quad (2.61)$$

We can see that there is an infinite magnification at $y = 0$, which is not physical. On the other hand, the total magnification tends to 1 when $y \rightarrow \infty$.

Such a model for a lens is not really realistic when considering lensing by galaxies or clusters of galaxies. Other models try to represent these better. Naturally, they are more complex, since they display a lower degree of symmetry and are more difficult to analyse than the simple model I have introduced. I will not present such models, but several different ones can be found in [22], [23] and [24].

Chapter 3

Gravitational lensing of gravitational waves: theoretical approach

In this chapter, we shall extend what we have seen with electromagnetic waves to gravitational waves. These will first be introduced and we will see what the differences with the electromagnetic case are. In particular, we will take into account wave effects and express an analytical solution to the problem of lensing of gravitational waves by a point-mass lens. This solution will then be discussed.

3.1 Introduction to gravitational waves

Before diving into the lensing of gravitational waves, one first needs to understand what they are. I will introduce these waves first conceptually, then mathematically. The wave equation that will be derived will also be used later in the study of lensed waves.

3.1.1 Gravitational waves

As explained in Section 1.1, matter deforms spacetime. Although the concept itself may be difficult to grasp, it would seem natural that, if it can be deformed, it can oscillate, and that these oscillations propagate through spacetime itself. These oscillations are *gravitational waves*. Their propagation is described by the wave equation, which we shall derive later. One can see such a wave as an oscillation of the curvature, hence of the metric. This means that the way to measure distances at a point changes because of the wave, but the coordinates of the point would not change.

Similarly to electric charges emitting radiation when undergoing an acceleration, massive objects emit gravitational radiations. However, contrarily to electromagnetic waves for which electric-dipole radiation dominates, gravitational waves are mostly quadrupolar, so that an acceleration is not sufficient, a change in the mass quadrupole moment (or a higher moment) is required. There exist multiple types of possible sources, among which we find black hole and neutron star mergers, pulsars and supernovae.

One characteristic of gravitational waves is that they tend to have a much lower frequency than visible light. For example, the frequency of the waves that are detected by current

interferometers is typically between 20 and 1000 Hz. This range is limited by the detectors and other waves should exist at much lower frequencies. The greater wavelength is intuitively natural, given the size of the systems that emit the wave. In the case of a binary black hole merger, the frequency of the wave is twice the frequency of the orbit [17]. It should also be noted that the amplitude of these waves is very small, typically 10^{-21} for a merger, and that, contrarily to light, they do not interact much with matter [25].

3.1.2 Wave equation

Now that we have conceptually introduced the subject, we will mathematically derive the expression of the gravitational wave, which is a solution to the *wave equation*. This equation expresses how a perturbation in the metric behaves. It is obtained by starting from Einstein's equations (1.16). As we shall see, the solution to the equation depends on the background metric. In other words, the presence of a lens that perturbs the background metric will produce a gravitational wave that is different from one where there is no lens. The complete derivation of the wave equation is quite cumbersome, so I will only present a few steps. The interested reader may however find a detailed development in Appendix B. The following developments are based on [26].

We consider a background metric $\tilde{g}_{\mu\nu}$, which is the metric describing the spacetime in which the wave propagates, and a perturbation in this metric $h_{\mu\nu}$, such that

$$g_{\mu\nu} = \tilde{g}_{\mu\nu} + h_{\mu\nu} . \quad (3.1)$$

Of course, we assume that the background metric is a solution to Einstein's equations (1.16). It is important to note that $h_{\mu\nu}$ is a perturbation, *i.e.* the background metric is such that it does not contain terms of the order of $h_{\mu\nu}$. In particular, the wavelength of the perturbation is much smaller than the characteristic scale at which the curvature of the background metric changes [27]. The idea to get the wave equation is to derive the *linearised Einstein's equations*, *i.e.* Einstein's equations for $h_{\mu\nu}$ to first order in $h_{\mu\nu}$ (meaning in the typical order of magnitude of the amplitude of $h_{\mu\nu}$). In general, typical amplitude values for $h_{\mu\nu}$ are around 10^{-21} . We will also consider that we solve the equation in vacuum.

So, to first order, we have

$$g^{\mu\nu} \simeq \tilde{g}^{\mu\nu} - h^{\mu\nu} , \quad (3.2)$$

since

$$\begin{aligned} \delta_{\mu}^{\nu} &= g_{\mu\alpha} g^{\nu\alpha} \\ &= (\tilde{g}_{\mu\alpha} + h_{\mu\alpha}) (\tilde{g}^{\nu\alpha} - h^{\nu\alpha}) \\ &= \delta_{\mu}^{\nu} - h_{\mu}^{\nu} + h_{\mu}^{\nu} - h_{\mu\alpha} h^{\nu\alpha} \\ &= \delta_{\mu}^{\nu} + O(h^2) , \end{aligned}$$

where $O(h^2)$ denotes second-order terms and h is used to ease the notation and represents here typical values of the elements of the $h_{\mu\nu}$ tensor. We then see that the equality is correct

to first order, so the inverse metric (3.2) is indeed the right one to first order. We can also notice that the tensor $\tilde{g}^{\mu\nu}$ raises indices and $\tilde{g}_{\mu\nu}$ lowers the indices of $h_{\mu\nu}$ to first order. To get the wave equation in a simple form, we make two main changes. The first is a change of function, by setting

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{h}{2} \tilde{g}_{\mu\nu} \quad \Rightarrow \quad \bar{h} = -h, \quad (3.3)$$

where we define $h = h_\alpha^\alpha$ the trace of the tensor. Therefore, the change of function reverses the trace. The second change is to use the De Donder gauge

$$\tilde{\nabla}_\alpha \bar{h}^{\alpha\beta} = 0, \quad (3.4)$$

i.e. change the coordinate system such that the latter relation holds¹. One can then show that Einstein's equations in vacuum to first order are

$$\tilde{G}_{\mu\nu} + \delta G_{\mu\nu} = 8\pi G \left(\tilde{T}_{\mu\nu} + \delta T_{\mu\nu} \right) \quad (3.5)$$

$$\Rightarrow \delta G_{\mu\nu} = \frac{1}{2} \left(-\tilde{\nabla}_\alpha \tilde{\nabla}^\alpha \bar{h}_{\mu\nu} + 2 \tilde{R}^\lambda{}_{\mu\nu\alpha} \bar{h}_\lambda{}^\alpha \right) = 8\pi G \delta T_{\mu\nu}, \quad (3.6)$$

where $\delta T_{\mu\nu}$ is related to the source of the gravitational wave, and $\delta G_{\mu\nu}$ is a linear perturbation (in h) of $G_{\mu\nu}$. If we assume that we are outside the source, we can take $\delta T_{\mu\nu} = 0$. We then obtain *linearised Einstein's equations in vacuum* and the full system, including the gauge condition, reads

$$\begin{cases} -\tilde{\nabla}_\alpha \tilde{\nabla}^\alpha \bar{h}_{\mu\nu} + 2 \tilde{R}^\lambda{}_{\mu\nu\alpha} \bar{h}_\lambda{}^\alpha = 0, \\ \tilde{\nabla}_\alpha \bar{h}^{\alpha\beta} = 0. \end{cases} \quad (3.7)$$

3.1.3 Solution for the Minkowski metric

We shall derive the solution of the wave equation in a Minkowski spacetime. This section is based on [28]. We can see that, in Minkowski spacetime where the connections and the Riemann tensor are null, the equations become a classical wave equation, with plane waves as solutions. We have

$$\begin{cases} \partial_\alpha \partial^\alpha \bar{h}_{\mu\nu} = 0, \\ \partial^\alpha \bar{h}_{\alpha\beta} = 0. \end{cases} \quad (3.8)$$

As the tensor $\bar{h}_{\mu\nu}$ is symmetric, it contains 10 independent elements, or functions (since $\bar{h}_{\mu\nu} = \bar{h}_{\nu\mu}(x)$). The gauge condition allows us to impose four constraints so that there remain 6 degrees of freedom. These can be reduced to 2 in several steps. First, we search for a plane wave solution that, without loss of generality, propagates along the z direction. As a result, $\bar{h}_{\mu\nu} = \bar{h}_{\mu\nu}(t-z)$ (plane wave solution) and thus $\partial_x \bar{h}_{\mu\nu} = \partial_y \bar{h}_{\mu\nu} = 0$. The gauge condition can now be written as

$$\partial_t \bar{h}_{t\beta} = \partial_z \bar{h}_{z\beta}, \quad (3.9)$$

¹It is possible to show that such a gauge always exists [17]

where $\bar{h}_{t\beta} \equiv \bar{h}_{0\beta}$ and $\bar{h}_{z\beta} \equiv \bar{h}_{3\beta}$. On the other hand, since $\bar{h}_{\mu\nu} = \bar{h}_{\mu\nu}(t - z)$, it satisfies

$$\partial_t \bar{h}_{\alpha\beta} - \partial_z \bar{h}_{\alpha\beta} = 0. \quad (3.10)$$

When we put these two results together, we find

$$\partial_t \bar{h}_{t\beta} = \partial_z \bar{h}_{z\beta} = 0, \quad (3.11)$$

so that

$$\bar{h}_{t\beta} = \bar{h}_{z\beta} = 0. \quad (3.12)$$

To be more precise, they are both equal to a constant that we set to 0. We can then represent the tensor as

$$\bar{h}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \bar{h}_{11} & \bar{h}_{12} & 0 \\ 0 & \bar{h}_{21} & \bar{h}_{22} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.13)$$

with $\bar{h}_{12} = \bar{h}_{21}$ since the tensor is symmetric. This leaves us with three degrees of freedom, the three independent elements. As said, we can reach two degrees of freedom. It must be noted that the trace is

$$\bar{h} = \bar{h}_{11} + \bar{h}_{22}, \quad (3.14)$$

and we can choose a frame of reference where this trace is null and the De Donder gauge is still valid. It is possible to show that such a reference frame exists. It is called the transverse-traceless gauge, or TT gauge. It is indeed transverse, since the time and longitudinal components are null and do not vary. In that gauge, $\bar{h} = h$ and $\bar{h}_{11} = -\bar{h}_{22}$. As a result, we can define

$$\bar{h}_+ = \bar{h}_{11} = -\bar{h}_{22} = h_+ \quad (3.15)$$

$$\bar{h}_\times = \bar{h}_{12} = \bar{h}_{21} = h_\times. \quad (3.16)$$

We can then finally express the solution in the TT gauge as

$$h_{\mu\nu}^{\text{TT}} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.17)$$

where h_+ is called the plus polarisation and h_\times the cross polarisation. This interpretation becomes more obvious when we write the interval

$$\begin{aligned} ds^2 = & -dt^2 + dz^2 + dx^2 (1 + h_+ \cos(\omega t - \omega z)) + dy^2 (1 - h_+ \cos(\omega t - \omega z)) \\ & + 2 dx dy (h_\times \cos(\omega t - \omega z)). \end{aligned} \quad (3.18)$$

Indeed, we see that the plus polarisation represents an oscillation of x and y , while the cross polarisation mixes the x and y components. These two polarisations are better understood with Fig. 3.1, where we see the $+$ wave oscillating in the shape of a '+', and the \times one as a cross. It also highlights the (dominant) quadrupolar nature of the oscillations.




Figure 3.1: Illustration of the $+$ and \times polarisations through the deformation of a ring of particles. Taken from [25].

3.2 Differences with the electromagnetic case

All the developments that were presented in Chapter 2 are valid in the geometrical optics approximation, *i.e.* we neglect all the wave effects such as diffraction. To be valid, it requires that the wavelength be much smaller than the characteristic dimension of the system, which is generally true in the context of lensing of electromagnetic waves (except maybe for radio waves). More formally, we usually consider that geometrical optics is valid when $\lambda \ll R_s$ [29], with $R_s = 2GM$ the Schwarzschild radius. Also, the geometrical optics description breaks down near caustics, since the magnification becomes infinite.

As stated earlier, gravitational waves are characterised by larger wavelengths, so that the condition $\lambda \ll R_s$ may not be fulfilled. For example, the frequency of 100 Hz roughly corresponds to a wavelength of about 3×10^6 m, so 3000 km. This scale is of the order of the characteristic size of some black holes and stars, more accurately of their Schwarzschild radius. Thus, in the case of microlensing, the approximation of geometrical optics does not hold anymore.

Nevertheless, in most cases, for example for strong lensing by a galaxy, the results of Chapter 2 are a good approximation, since in that case $\lambda \ll R_s$. Also, most of the things we shall discuss are also valid for electromagnetic radiation, since the wave equation with a lens that we will analyse does not take into account the polarisation, but only a scalar field. The same approximation can then be applied to electromagnetic radiation, and the equations for gravitational waves and electromagnetic ones will be the same.

3.3 Taking wave effects into account

As already stated, the developments of Chapter 2 were made in the geometrical optics approximation, but the wave nature of light can also be taken into account. What follows is also valid for gravitational waves. This section is largely based on [24].

When one solves the equation for an electromagnetic field propagating in a metric perturbed by a lens using the Fresnel-Kirchhoff diffraction integral, one finds, for a point source,

$$\vec{\tilde{E}}(\omega) = A V(\omega) \vec{C}(\omega), \quad (3.19)$$

where $\vec{\tilde{E}}(\omega)$ is the Fourier transform of the electric field, A is a constant, $V(\omega)$ is called the transmission factor and contains all the influence of the lens on the wave, and $\vec{C}(\omega)$ is related to the emission process. In principle, both V and \vec{C} depend on $\vec{\eta}$, the position of the source, but it is fixed in this case, so I will not write the dependence explicitly. If the source is extended, one needs to perform an integral over the source plane.

As said, V accounts for the effect of the lens and is thus the main difference with the unlensed case. Following [30], I will also refer to it as the *amplification factor* $F(\omega)$, which is the ratio of the Fourier transform coefficients of the lensed over unlensed waves. The transmission factor is defined as

$$V(\omega) = \int_{\mathbb{R}^2} e^{i\omega \Delta t(\vec{x})} d\vec{x} = \bar{F}(\omega), \quad (3.20)$$

where $\Delta t(\vec{x})$ is the time delay as defined in Eq. (2.24), and \bar{F} is the unnormalised amplification factor. In other words, it takes into account the contributions of the waves that come from the whole lens plane, due to wave effects. However, the main contributions will come from the regions where the images are, *i.e.* where Δt is stationary. In fact, the frequency of the exponential is very large for large ω (*i.e.* in the geometrical optics approximation), since the multiplicative constants of the time delay are very large. It appears that, at low frequencies, the contributions from the whole lens plane matter, while at high frequencies (where geometrical optics applies) the contributions from other parts than the stationary points vanish because of destructive interference [29]. We can introduce

$$\tilde{\omega} = \omega \frac{D_S}{D_L D_{LS}} (1 + z_L) \xi_0^2 = \omega \frac{D_S D_L}{D_{LS}} (1 + z_L), \quad (3.21)$$

and the amplification factor becomes

$$\bar{F}(\omega) = \int_{\mathbb{R}^2} e^{i\tilde{\omega} \Psi(\vec{x})} d\vec{x}, \quad (3.22)$$

where $\Psi(\vec{x})$ is Fermat's potential (2.25). The fact that $\tilde{\omega}$ is very large allows us to use an asymptotic expansion of the integral. This can be done using the method of stationary phase, which is used to expand the integral around stationary points, *i.e.* the images. A general expression of this method is given in 1D in [31]. It reads:

$$I = \int_{\alpha}^{\beta} e^{i\nu f(x)} \varphi(x) dx = \left(\frac{\pi}{\epsilon 2\nu f''(\alpha)} \right)^{1/2} \varphi(\alpha) e^{i\nu f(\alpha) + \epsilon i\pi/4} + O\left(\frac{1}{\nu}\right), \quad \nu \rightarrow \infty, \quad (3.23)$$

where $f''(x)$ is the second derivative of f , $\epsilon = 1$ if $f''(\alpha) > 0$ and $\epsilon = -1$ if $f''(\alpha) < 0$. In our case, $\phi(x) = 1$, $\nu = \tilde{\omega}$ and $f(x) = \Psi(\vec{x})$. This method can be generalised to higher dimensions, which is what we need. The form in two dimensions is given, for example, in [32]. If the stationary point of f is (α_1, α_2) , the integral is expressed as

$$I = \frac{2i\pi\sigma}{\nu} |\det A^*|^{-1/2} \varphi(\alpha_1, \alpha_2) e^{i\nu f(\alpha_1, \alpha_2)}, \quad (3.24)$$

where A^* is the Hessian matrix evaluated at the stationary points, which intuitively arises from the expansion around (α_1, α_2) and generalises the second derivative of the one-dimensional case. The parameter σ is such that [32]

$$\sigma = \begin{cases} 1 & \text{if } \det A^* > 0 \text{ and } A_{00}^* > 0, \\ -1 & \text{if } \det A^* > 0 \text{ and } A_{00}^* < 0, \\ -i & \text{if } \det A^* < 0. \end{cases} \quad (3.25)$$

In our case, one expands Fermat's potential around the images, which are stationary points. Similarly to A^* , I shall refer to the quantity evaluated at the stationary points with a star, *i.e.* Ψ^* is Fermat's potential at the image position. The Hessian matrix A^* corresponds to the Jacobian matrix A defined in Eq. (2.42), with which the types of images are defined (see Section 2.6). We can then apply the method of stationary phases to our problem to write, for one image,

$$\bar{F}(\omega) = \frac{2i\pi\sigma}{\tilde{\omega}} |\det A^*|^{-1/2} e^{i\tilde{\omega}\Psi^*}. \quad (3.26)$$

We can also notice that the definition (3.25) exactly corresponds to the different types of images, so that we can introduce

$$\sigma = e^{-ni\pi/2}, \quad (3.27)$$

with n the *Morse index*, being equal to 0, 1, 2 for type I, II and III images respectively. It corresponds to the number of caustics crossed by the ray when travelling towards the observer. We have seen that caustics are defined as curves where $\det A = 0$. So, intuitively, crossing a caustic would change the sign of the determinant. Also, all images near the source are of type I (since there is no lensing). Therefore, type I images would not cross any caustic, type II ones, for which $\det A < 0$, would cross once a caustic, while type III ones ($\det A > 0$, $\text{Tr } A < 0$) would cross it twice to change twice the sign of the determinant. The phase represented by the index n is referred to as the *Morse phase* [12]. We can also notice that $|\det A^*|^{-1}$ is the absolute value of the magnification μ .

Finally, we need to account for all the images. The integral in the definition of \bar{F} can be split into a sum of integrals, with one around each image on an arbitrarily small disc, and the others on the rest of the plane. The latter integrals average out to zero because of the large integration domain and high frequency of the exponential, while the integrals over the images give each the expression (3.26). As a result, we get a sum over all the images, *i.e.*

$$\bar{F}(\omega) = \frac{2i\pi}{\omega} \frac{D_{LS}}{D_S D_L} \frac{1}{1+z_L} \sum_j^N |\mu_j|^{1/2} e^{i\omega \Delta t_j^* - n_j i\pi/2}, \quad (3.28)$$

where N is the number of images and μ_j the magnification of the j th image and n_j the Morse index of the j th image. Given the presence of the inverse of $\det A$, this development is not valid near caustics.

To end this section, we introduce the true amplification factor. We need to normalise $V = \bar{F}$, so that $|F| = 1$ when there is no lensing. We therefore need to compute \bar{F} for $\Psi(\vec{x}) = \frac{1}{2}(\vec{x} - \vec{y})^2$ (no lensing), with $\vec{y} = 0$ (by choice of the reference axis), so we get

$$\bar{F}(\omega) = \int_{\mathbb{R}^2} e^{i\tilde{\omega}(x_1^2+x_2^2)/2} d\vec{x} = \frac{2i\pi}{\tilde{\omega}}. \quad (3.29)$$

We thus express the amplification factor as

$$F(\omega) = \frac{\tilde{\omega}}{2i\pi} \bar{F}(\omega) = \sum_j^N |\mu_j|^{1/2} e^{i\omega \Delta t_j^* - n_j i\pi/2}. \quad (3.30)$$

We can also compute its modulus in the case of a point-mass lens, which will be useful later. with such a lens, we have two images, hence

$$\begin{aligned} |F(\omega)| &= \left| |\mu_+|^{1/2} e^{i\omega \Delta t_+^* - i(\pi/2) n_+} + |\mu_-|^{1/2} e^{i\omega \Delta t_-^* - i(\pi/2) n_-} \right| \\ &= \sqrt{|\mu_+| + |\mu_-| + |\mu_+|^{1/2} |\mu_-|^{1/2} (e^{i\omega \Delta t_d - i \Delta n \pi/2} + e^{-i\omega \Delta t_d + i \Delta n \pi/2})} \\ &= \sqrt{|\mu_+| + |\mu_-| + 2 |\mu_+|^{1/2} |\mu_-|^{1/2} \cos(\omega \Delta t_d - \Delta n \pi/2)}, \end{aligned} \quad (3.31)$$

where Δt_d is the time delay between the two images, and Δn the difference in the Morse indices. Given the definition of the cosine, we can equivalently use $\Delta t_d = \Delta t_+ - \Delta t_-$ with $\Delta n = n_+ - n_-$, or $\Delta t_d = \Delta t_- - \Delta t_+$ with $\Delta n = n_- - n_+$. We have seen in Section 2.7 that μ_+ is always greater than 1, so that the $+$ image is of type I or III, so n_+ is 0 or 2. On the other hand, $\mu_- \leq 0$, so the $-$ image is of type II, hence $n_- = 1$. As a result, $\Delta n = \pm 1$, but we can define it in such a way that it is always 1 by choosing the order of the indices properly (*i.e.* n_+ or n_- first). The order of the indices in Δt_d will be the same as the one chosen for Δn . We can then rewrite the cosine as $\cos(\omega \Delta t_d - \pi/2) = \sin(\omega \Delta t_d)$, and the modulus of the amplification factor becomes

$$|F(\omega)| = \sqrt{|\mu_+| + |\mu_-| + 2 |\mu_+|^{1/2} |\mu_-|^{1/2} \sin(\omega \Delta t_d)}. \quad (3.32)$$

The amplification factor expresses how the wave is modified in both amplitude and phase. We also see that, contrarily to what was done in Chapter 2, we compute how the image is modified according to its frequency. It is important to keep in mind that the result (3.30) is valid where geometrical optics applies. To take into account wave effects fully, one would need to solve the integral in the general case where ω is not so large. This is what will be done in the following sections, where we will solve the wave equation with the presence of a lens, and the solution will be the same as the one obtained by the calculation of Eq. (3.22).

3.4 Wave equation with a point-mass lens

We will now investigate mathematically what happens to the wave in the presence of a point-mass lens. To do so, we will use the wave equation that we have derived. To ease the notation, I will drop all the bars and tildes, as they are not really important anymore in what follows. We consider the background metric as being the metric perturbed by a mass, the lens, such that its gravitational potential $U(\vec{r})$ is small, *i.e.* $U(\vec{r}) \ll 1$. The potential is

$$U(\vec{r}) = -\frac{GM}{|\vec{r}|}, \quad (3.33)$$

and

$$\vec{\nabla}U(\vec{r}) = \frac{GM \vec{r}}{|\vec{r}|^3} \quad \Rightarrow \quad \left| \vec{\nabla}U(\vec{r}) \right| = \frac{GM}{|\vec{r}|^2}. \quad (3.34)$$

We then see that $\left| \vec{\nabla}U \right|$ is smaller than U , since we consider that $|\vec{r}| \gg GM$ (as $U \ll 1$). We thus have, in the reference frame of the lens,

$$ds^2 = -(1 + 2U(\vec{r})) dt^2 + (1 - 2U(\vec{r})) d\vec{r}^2,$$

which represents the metric (2.11), *i.e.* the one of a Minkowski spacetime perturbed by a weak lens. In other words, the metric is

$$g_{\mu\nu} = \text{diag}(-1 - 2U, 1 - 2U, 1 - 2U, 1 - 2U), \quad (3.35)$$

where $U = U(\vec{r})$ to ease the notation. We can then deduce that

$$g^{\mu\nu} = \text{diag}\left(\frac{1}{-1 - 2U}, \frac{1}{1 - 2U}, \frac{1}{1 - 2U}, \frac{1}{1 - 2U}\right). \quad (3.36)$$

We can also write the determinant of the metric:

$$g = \det(g_{\mu\nu}) = -(1 + 2U)(1 - 2U)^3, \quad (3.37)$$

which is always negative since U is very small.

We will assume that the wavelength is much smaller than the typical radius of curvature of the background metric, so that Eq. (3.7) reduces to [30]

$$\nabla_\alpha \nabla^\alpha h_{\mu\nu} \equiv \square h_{\mu\nu} = 0, \quad (3.38)$$

with \square the d'Alembertian built on the background metric. We will further assume that the polarisation of the wave remains unchanged after lensing [30], so that we may express the equation only for $\Phi(\vec{r}, t)$, if $h_{\mu\nu} = \Phi e_{\mu\nu}$ with $e_{\mu\nu}$ the polarisation tensor. Again, I will write $\Phi \equiv \Phi(\vec{r}, t)$ to ease the notation. In this case, it can be shown (see Appendix C) that the equation can be rewritten as

$$\partial_\mu (\sqrt{-g} g^{\mu\nu} \partial_\nu \Phi) = 0, \quad (3.39)$$

which is more convenient to use, since it does not require to compute the connection coefficients. We then start from the wave equation for the amplitude:

$$\partial_\mu (\sqrt{-g} g^{\mu\nu} \partial_\nu \Phi) = \partial_\mu ((1 + 2U)(1 - 2U)^3 g^{\mu\nu} \partial_\nu \Phi) = 0. \quad (3.40)$$

There are quite a few terms, so we can consider them one by one. The potential does not depend on time, which simplifies the temporal term to:

$$\begin{aligned} \partial_0 \left[\sqrt{(1+2U)(1-2U)^3} g^{0\nu} \partial_\nu \Phi \right] &= \sqrt{(1+2U)(1-2U)^3} g^{00} \partial_0 \partial_0 \Phi \\ &= -\sqrt{1-4U^2} \frac{1-2U}{1+2U} \partial_0 \partial_0 \Phi. \end{aligned} \quad (3.41)$$

Since the spatial terms of the metric are the same, we can write them together with

$$\begin{aligned} \partial_i \left[\sqrt{(1+2U)(1-2U)^3} g^{i\nu} \partial_\nu \Phi \right] &= \partial_i \left[\sqrt{(1+2U)(1-2U)^3} \frac{1}{1-2U} \partial_i \Phi \right] \\ &= \partial_i \left[\sqrt{1-4U^2} \partial_i \Phi \right]. \end{aligned} \quad (3.42)$$

The derivative of the potential term is

$$\partial_i \left(\sqrt{1-4U^2} \right) = -\frac{4U}{\sqrt{1-4U^2}} \partial_i U, \quad (3.43)$$

so that the spatial terms can be written as

$$\partial_i \left[\sqrt{(1+2U)(1-2U)^3} g^{i\nu} \partial_\nu \Phi \right] = -\frac{4U}{\sqrt{1-4U^2}} \partial_i U \partial_i \Phi + \sqrt{1-4U^2} \partial_i \partial_i \Phi \quad (3.44)$$

We then put together equations (3.40), (3.41) and (3.44) and divide by $\sqrt{1-4U^2}$ to get:

$$-\frac{1-2U}{1+2U} \partial_0 \partial_0 \Phi - \frac{4U}{1-4U^2} \partial_i U \partial_i \Phi + \partial_i \partial_i \Phi = 0, \quad (3.45)$$

where the sum is implicit, or in more familiar notations

$$-\frac{1-2U}{1+2U} \partial_0^2 \Phi - \frac{4U}{1-4U^2} \vec{\nabla} U \cdot \vec{\nabla} \Phi + \vec{\nabla}^2 \Phi = 0, \quad (3.46)$$

where all the operators now refer to Minkowski operators in Cartesian coordinates, *i.e.* taking the identity matrix as spatial metric (so $\vec{\nabla}^2 = \sum_i \partial_i^2$ and $\vec{a} \cdot \vec{b} = \sum_i a_i b_i$). We have that

$$\frac{4U}{1-4U^2} = 4U + O(U^2), \quad (3.47)$$

so, since $\partial_i U$ is $O(U)$ (see (3.34)) and if we then assume that $\partial_i \phi$ is $O(1)$, the second term of (3.45) is $O(U^2)$. One can also easily show that

$$\frac{1-2U}{1+2U} = 1 - 4U + O(U^2). \quad (3.48)$$

So, to first order in U , the wave equation becomes

$$\vec{\nabla}^2 \Phi - (1-4U) \partial_0^2 \Phi = 0. \quad (3.49)$$

This equation is thus the wave equation that we have to solve in our case.

3.5 Analytical solution for a point-mass lens

In this section, we will solve the equation for a plane wave propagating in the metric perturbed by the point-mass lens. The solution will be illustrated and commented as well.

3.5.1 Derivation of the solution

First, let us clarify the problem we are trying to solve. We consider a plane wave propagating towards a point-mass lens in a metric perturbed by the latter. We are interested in finding how the lens will modify the initial wave from the point of view of the observer. To this end, we will solve the wave equation (3.49), which is expressed in the reference frame of the lens and describes how the metric influences the propagation of the wave.

We assume that the solution can be written in the form $\Phi(\vec{r}, t) = \phi(\vec{r}) \exp(-i\omega t)$. We could also take into account the phase of the wave, *i.e.* $\Phi(\vec{r}, t) = \phi(\vec{r}) \exp(-i\omega(t - t_0))$. We shall however not consider it in the following, without loss of generality. We can then write Eq. (3.49) as

$$\vec{\nabla}^2 \Phi + i\omega(1 - 4U) \partial_0 \Phi = 0, \quad (3.50)$$

given the properties of the exponential. We can further develop this equation to get

$$\begin{aligned} 0 &= \vec{\nabla}^2 \Phi + i\omega(1 - 4U) \partial_0 \Phi \\ &= \vec{\nabla}^2 \Phi + i\omega \partial_0 \Phi - 4U \omega^2 \Phi \\ \Leftrightarrow i\partial_0 \Phi &= -\frac{1}{\omega} \vec{\nabla}^2 \Phi + V \Phi, \end{aligned} \quad (3.51)$$

with $V = 4U\omega$. The last equation was formulated in such a way that it makes a famous equation appear. Indeed, since we can consider here ω as a constant parameter, we retrieve a Schrödinger-like equation! Fortunately, the analytical solution to such an equation is known, at least for a few V . We will rewrite the gravitational potential as

$$U = -\frac{GM}{r} = -\frac{R_s}{2r}, \quad (3.52)$$

with $R_s = 2GM$ the Schwarzschild radius. In this case, solving our ‘Schrödinger equation’ is actually similar to solving the Schrödinger equation in a Coulomb potential! To be more accurate, it is similar to the elastic scattering in a Coulomb potential, and the solution to that problem is known. I have taken the answer to the electromagnetic problem from [33] and I will show it here adapted to our problem. Such a method to solve the equation was already proposed by, among others, [34] and [35].

We can further develop Eq. (3.51):

$$\begin{aligned} i\partial_0 \Phi &= -\frac{1}{\omega} \vec{\nabla}^2 \Phi + V \Phi \\ \Rightarrow 0 &= -\vec{\nabla}^2 \phi - \omega^2 \phi + 4U\omega^2 \phi \end{aligned} \quad (3.53)$$

$$\Rightarrow 0 = \vec{\nabla}^2 \phi + \omega^2 \left(1 + \frac{2R_s}{r}\right) \phi. \quad (3.54)$$

We will now use parabolic coordinates, given the axis symmetry of the system. We define the z -axis as being parallel to the direction of the wave vector of the incoming wave. We then perform the following change of variables:

$$\xi = r - z, \quad (3.55)$$

$$\eta = r + z. \quad (3.56)$$

To be more precise, it corresponds to the following change of variables:

$$x = \sqrt{\xi\eta} \cos \varphi, \quad (3.57)$$

$$y = \sqrt{\xi\eta} \sin \varphi, \quad (3.58)$$

$$z = \frac{1}{2}(\eta - \xi), \quad (3.59)$$

with then $r = \sqrt{x^2 + y^2 + z^2}$ and $\tan \varphi = y/x$. In these coordinates, the Laplace operator becomes [33]

$$\vec{\nabla}^2 = \frac{4}{\xi + \eta} \left[\frac{\partial}{\partial \xi} \left(\xi \frac{\partial}{\partial \xi} \right) + \frac{\partial}{\partial \eta} \left(\eta \frac{\partial}{\partial \eta} \right) \right] + \frac{1}{\xi\eta} \frac{\partial^2}{\partial \varphi^2}. \quad (3.60)$$

Given the symmetry of the problem, the last derivative is null, *i.e.* $\phi(\vec{r}) = \phi(\xi, \eta)$.

If we apply the change of coordinates to Eq. (3.54), with $2r = \xi + \eta$, we get

$$\frac{4}{\xi + \eta} \left[\frac{\partial}{\partial \xi} \left(\xi \frac{\partial \phi}{\partial \xi} \right) + \frac{\partial}{\partial \eta} \left(\eta \frac{\partial \phi}{\partial \eta} \right) \right] + \omega^2 \left(1 + \frac{4R_s}{\xi + \eta} \right) \phi = 0. \quad (3.61)$$

We then set

$$\phi(\xi, \eta) = \exp \left[\frac{i}{2} \omega (\eta - \xi) \right] \psi(\xi). \quad (3.62)$$

We can see that the exponential corresponds to a wave propagating along z , while its amplitude depends on its position.

We then inject this expression into Eq. (3.61), but this yields a lot of terms. We will thus express the Laplace operator in several steps to ease the notation. We can first notice that the exponential will be present in each term in the end, since the derivatives do not modify it, and can thus be cancelled out. We also have that

$$\frac{\partial \phi}{\partial \xi} = \left(\frac{\partial \psi}{\partial \xi} - \frac{i}{2} \omega \psi \right) \exp \left[\frac{i}{2} \omega (\eta - \xi) \right],$$

$$\frac{\partial \phi}{\partial \eta} = \frac{i}{2} \omega \psi \exp \left[\frac{i}{2} \omega (\eta - \xi) \right].$$

For the terms in the Laplace operator, we thus get

$$\frac{\partial}{\partial \xi} \left(\xi \frac{\partial \phi}{\partial \xi} \right) = \left\{ \frac{\partial \psi}{\partial \xi} - \frac{i}{2} \omega \psi + \xi \left[\frac{\partial^2 \psi}{\partial \xi^2} - \frac{i\omega}{2} \frac{\partial \psi}{\partial \xi} - \frac{i}{2} \omega \left(\frac{\partial \psi}{\partial \xi} - \frac{i}{2} \omega \psi \right) \right] \right\} e^{\frac{i}{2} \omega (\eta - \xi)}, \quad (3.63)$$

$$\frac{\partial}{\partial \eta} \left(\eta \frac{\partial \phi}{\partial \eta} \right) = \left\{ \frac{i}{2} \omega \psi - \frac{1}{4} \omega^2 \eta \psi \right\} e^{\frac{i}{2} \omega (\eta - \xi)}. \quad (3.64)$$

We can see that the second term of the first equality cancels the first term of the second equality. Simplifying these expressions and cancelling out the exponential, Eq. (3.61) becomes

$$\frac{4}{\xi + \eta} \left[\xi \frac{\partial^2 \psi}{\partial \xi^2} + (1 - i\omega\xi) \frac{\partial \psi}{\partial \xi} - \frac{1}{4} \omega^2 (\xi + \eta) \psi \right] + \omega^2 \left(1 + \frac{4R_s}{\xi + \eta} \right) \psi = 0. \quad (3.65)$$

We can see that the third term cancels with the fourth one. We then finally end up with

$$\xi \frac{d^2 \psi}{d\xi^2} + (1 - i\omega\xi) \frac{d\psi}{d\xi} + \omega^2 R_s \psi = 0, \quad (3.66)$$

where the partial derivatives were replaced by total ones, since ψ only depends on ξ .

So, we have arrived at an equation for ψ . It appears that this equation has a well-known form, which is Kummer's equation [36]

$$\hat{z} \frac{d^2 w}{d\hat{z}^2} + (b - \hat{z}) \frac{dw}{d\hat{z}} - a w = 0, \quad (3.67)$$

with $\hat{z} \in \mathbb{C}$. This equation has a solution

$${}_1F_1(a, b, \hat{z}) = \sum_{n=0}^{\infty} \frac{(a)_n}{(b)_n n!} \hat{z}^n, \quad (3.68)$$

with the $(a)_n$ the Pochhammer's symbol that denotes the rising factorial², *i.e.* [36]

$$(a)_n = \begin{cases} (a)(a+1)(a+2)\dots(a+n-1) & \text{if } n > 0, \\ 1 & \text{if } n = 0, \end{cases} \quad (3.69)$$

or in general [36]

$$(a)_n = \frac{\Gamma(a+n)}{\Gamma(a)} \quad a \neq 0, -1, -2, \dots, \quad (3.70)$$

with $\Gamma(\hat{z})$ the gamma function. The function ${}_1F_1$ is called the confluent hypergeometric function, or also Kummer's function (noted $M(a, b, \hat{z})$).

From Kummer's equation (3.67), we need to set $\hat{z} = i\omega\xi$ to retrieve our equation (3.66). We then indeed obtain

$$\begin{aligned} i\omega\xi \frac{1}{(i\omega)^2} \frac{d^2 w}{d\xi^2} + (b - i\omega\xi) \frac{1}{i\omega} \frac{dw}{d\xi} - a w &= 0 \\ \Rightarrow \xi \frac{d^2 w}{d\xi^2} + (b - i\omega\xi) \frac{dw}{d\xi} - i\omega a w &= 0 \end{aligned} \quad (3.71)$$

We can compare it to equation (3.66) and we find that $b = 1$ and $a = i\omega R_s$, so that the solution is

$$\psi(\xi) = C {}_1F_1(i\omega R_s, 1, i\omega\xi), \quad (3.72)$$

²In combinatorics, the rising factorial is noted $a^{\bar{n}}$.

with C a constant. The full solution to the problem in parabolic coordinates is

$$\Phi(\xi, \eta, t) = C {}_1F_1(i\omega R_s, 1, i\omega\xi) \exp\left[\frac{i}{2}\omega(\eta - \xi)\right] \exp(-i\omega t). \quad (3.73)$$

For convenience, we may re-express the solution in spherical coordinates. One way to do it is to notice that $\xi = 2r \sin^2(\theta/2)$. Indeed,

$$\xi = r - z = r - r \cos \theta = r(1 - \cos \theta) = 2r \sin^2\left(\frac{\theta}{2}\right). \quad (3.74)$$

The other terms can easily be replaced, and we thus get

$$\Phi(\vec{r}, t) = C {}_1F_1\left(i\omega R_s, 1, 2i\omega r \sin^2\left(\frac{\theta}{2}\right)\right) \exp(i\omega z) \exp(-i\omega t), \quad (3.75)$$

or with $z = r \cos \theta$,

$$\Phi(\vec{r}, t) = C {}_1F_1\left(i\omega R_s, 1, 2i\omega r \sin^2\left(\frac{\theta}{2}\right)\right) \exp(i\omega r \cos \theta) \exp(-i\omega t). \quad (3.76)$$

Note that in SI units, we would have all the ω that are divided by c in ${}_1F_1$ and in the first exponential. This solution is the same as the one obtained by [34]. It is also expressed for a monochromatic wave. In the general case, we write

$$\Phi(\vec{r}, t) = \int_{-\infty}^{+\infty} \phi(\vec{r}, \omega) e^{-i\omega t} \frac{d\omega}{2\pi}. \quad (3.77)$$

We then directly retrieve Eq. (3.54) for ϕ and its solution is exactly the same. The only difference is that we must integrate Eq. (3.76) over ω . In the following, we shall consider the general case with $\phi(\vec{r}, \omega)$.

To extract what is really interesting, we can express the amplification factor F , which is the ratio of ϕ_L (lensed wave) over ϕ (initial wave), as introduced in Section 3.3. In this case, one would have

$$F(\vec{r}, \omega) = C {}_1F_1\left(i\omega R_s, 1, 2i\omega r \sin^2\left(\frac{\theta}{2}\right)\right), \quad (3.78)$$

since the exponential is the same for both waves.

So, we have found an analytical solution, but its complicated form, particularly with the confluent hypergeometric function, does not allow for an easy interpretation. We shall therefore investigate if we can express the solution differently and then interpret the results.

3.5.2 Re-expressing the solution

A first important remark is that we do not need to change the coordinate system, centred on the lens, in order to describe what the observer would observe. In fact, it would only make the solution less elegant by introducing new terms. We consider the same configuration as in Chapter 2, but the way we tackle the problem is different. Indeed, as shown in Fig. 3.2a, we do not consider a *ray* anymore, but really consider the *wave* and seek its deformation by

the lens. The general geometry, shown in Fig. 3.2b, is still the same. We need to express the solution at the coordinates of the observer, *i.e.* $r = D_L$ and $\theta = \gamma$, where γ is the angle between the initial wave vector, which defines the z -axis, and the reference axis. Note that the distances are not angular diameter ones, since we do not take expansion into account. It is however not important, as the distances will not appear in the final expression of the solution. Since the position is fixed, we will see how F behaves as a function of ω only. The




Figure 3.2: System of coordinates centred on the lens. P is a point at which we evaluate the wave. The angles have been largely exaggerated. (a) Zoom in on the lens to define the coordinate system with the incoming wave propagating at an angle γ with respect to the reference axis. (b) View of the angle γ in the whole system (if the source is not at an infinite distance).

amplification factor at O is thus

$$F(\omega) = C {}_1F_1 \left(i\omega R_s, 1, 2i\omega D_L \sin^2 \left(\frac{\gamma}{2} \right) \right). \quad (3.79)$$

From Fig. 3.2b, it is straightforward to get $\gamma = D_S \beta / D_{LS}$ (given the small-angle approximation). However, to have a plane wave coming on the lens, one needs to consider the source as being infinitely distant. Otherwise, the wave is spherical and the curvature of the wavefronts must be taken into account. As a consequence, it does not really make sense to take into account the infinite distances D_S and D_{LS} . We can see that, in the limit $D_S \rightarrow \infty$, $\gamma \simeq \beta$. Since we consider a point-mass lens, we use $y = \beta / \theta_E$ (see Section 2.7), with θ_E the Einstein radius introduced in Eq. (2.54). Note that it induces the characteristic length $\xi_0 = D_L \theta_E$. Given the infinite distance to the source, we need to define a more appropriate characteristic angle. We will then rather use [35]

$$\theta_E^\infty = \sqrt{\frac{2R_s}{D_L}}, \quad (3.80)$$

where the ∞ symbol refers to the fact that the source is at an infinite distance. In that context, we set $y = \beta / \theta_E^\infty$, which is then interpreted as the dimensionless angular position of the source. Moreover, given the small-angle approximation, we have that $\sin^2(x) \simeq x^2$. As a result, we can replace γ and inject the definition of θ_E^∞ into the expression of the amplification factor to get

$$F(\omega) = C {}_1F_1 \left(i\omega R_s, 1, 2i\omega D_L \frac{y^2 \theta_E^{\infty 2}}{4} \right) = C {}_1F_1 \left(i\omega R_s, 1, i\omega R_s y^2 \right). \quad (3.81)$$

To simplify the notation of this expression further, we introduce $\tilde{\omega} = \omega R_s$, which is equivalent (within a factor 2) to the one defined in (3.21), but in a non-expanding Universe and with $\xi_0 = D_L \theta_E$. The amplification factor then becomes

$$F(\omega) = C {}_1F_1(i\tilde{\omega}, 1, i\tilde{\omega} y^2) . \quad (3.82)$$

We still have to determine the constant C to normalise the amplification factor properly. What can be done is to normalise it so that its modulus is 1 at infinity [37]. To this end, we can use the asymptotic expansion of the confluent hypergeometric function for a large argument [36]:

$${}_1F_1(a, b, \hat{z}) \sim \frac{\Gamma(b)}{\Gamma(a)} e^{\hat{z}} \hat{z}^{a-b} \sum_{s=0}^{\infty} \frac{(1-a)_s (b-a)_s}{s!} \hat{z}^{-s} + \frac{\Gamma(b) (-\hat{z})^{-a}}{\Gamma(b-a)} \sum_{s=0}^{\infty} \frac{(a)_s (a-b+1)_s}{s!} (-\hat{z})^{-s} . \quad (3.83)$$

when $|\hat{z}|$ is large. In our case, we look for $|\hat{z}| \rightarrow \infty$, with $a = i\tilde{\omega}$ and $b = 1$. There is only one term that will then not vanish, since the first term is always $O(\hat{z}^{-1})$ because of the \hat{z}^{-b} factor. We thus get

$${}_1F_1(i\tilde{\omega}, 1, \hat{z}) \sim \frac{1}{\Gamma(1-i\tilde{\omega})} (-\hat{z})^{-i\tilde{\omega}}, \quad \hat{z} \rightarrow \infty . \quad (3.84)$$

As $\hat{z} = 2i\omega r \sin^2(\theta/2) = i\omega\xi$ (to ease the notation), we have

$$(-\hat{z})^{-i\tilde{\omega}} = (-i\omega\xi)^{-i\tilde{\omega}} = \exp\left[\ln\left((-i\omega\xi)^{-i\tilde{\omega}}\right)\right] = \exp[-i\tilde{\omega} \ln(-i) - i\tilde{\omega} \ln(\omega\xi)] = e^{-\tilde{\omega}\pi/2} e^{-i\tilde{\omega} \ln(\omega\xi)} \quad (3.85)$$

We can inject this result into Eq. (3.84) to finally get

$${}_1F_1(i\tilde{\omega}, 1, \hat{z}) \sim \frac{e^{-\tilde{\omega}\pi/2}}{\Gamma(1-i\tilde{\omega})} e^{-i\tilde{\omega} \ln(\omega\xi)}, \quad \hat{z} \rightarrow \infty . \quad (3.86)$$

For $|F| = 1$ at infinity, we thus set

$$C = \Gamma(1-i\tilde{\omega}) e^{\tilde{\omega}\pi/2} e^{i\alpha} , \quad (3.87)$$

where $\alpha \in \mathbb{R}$. We can indeed add a phase term that will not change the modulus of the amplification factor. The choice of α is not of great importance and one can set it to 0. It can equivalently be absorbed in the phase of the wave. The final form of the solution is thus

$$F(\omega) = \Gamma(1-i\tilde{\omega}) e^{\tilde{\omega}\pi/2} {}_1F_1(i\tilde{\omega}, 1, i\tilde{\omega} y^2) . \quad (3.88)$$

It can also be found under different forms in different papers (among others [29], [35] and [30]), where α may have been chosen differently. It can also be found with a $\tilde{\omega} \rightarrow \tilde{\omega}/2$ by considering $\tilde{\omega} = \omega 4GM = 2\omega R_s$, while I have defined $\tilde{\omega} = \omega R_s$.

It is interesting to note that this solution can also be found from the Fresnel-Kirchhoff diffraction integral method, which led to the definition of the transmission factor V (see Eq. (3.20)). It is indeed possible to directly integrate the time delay (with $\xi_0 = R_E$) of a point source ($\propto (\vec{x}^2 - \vec{y}^2)/2 - \ln(|\vec{x}|)$). For more details, see [24].

3.5.3 Interpreting the solution

The asymptotic expansion (3.83) of the confluent hypergeometric function can be useful for interpreting the solution. Limiting ourselves to the first order and using $\Gamma(\hat{z}) = \Gamma(\hat{z}+1)/\hat{z}$, we have

$$\begin{aligned} \phi(r, \theta, \omega) \simeq & e^{-i\tilde{\omega} \ln(2\omega r \sin^2(\theta/2))} \left[1 - \frac{\tilde{\omega}^2}{2\omega r \sin^2(\theta/2)} \right] e^{i\omega r \cos(\theta)} \\ & + \frac{\Gamma(1 - i\tilde{\omega})}{\Gamma(1 + i\tilde{\omega})} \frac{\tilde{\omega}}{2\omega \sin^2(\theta/2)} \frac{e^{i\tilde{\omega} \ln(2\omega r \sin^2(\theta/2))}}{r} e^{i\omega r}, \end{aligned} \quad (3.89)$$

so that the resulting wave is a superposition of a (distorted) plane wave and an outgoing spherical wave. However, in the case of gravitational lensing, the angles are small so that $r \sin^2(\theta/2)$ may not be that large and the asymptotic expansion would then not be appropriate [35]. The best method to try and interpret the solution is then to compute it numerically and visualise the result.

We will focus on the interpretation of the solution at the observer, *i.e.* Eq. (3.88). We can look at how the amplification factor amplitude changes with both y and ω . This is represented in Fig. 3.3. First, we can see in Fig. 3.3a that the global evolution of $|F|$ with y is the same at all frequencies, *i.e.* it oscillates around 1 and is maximum at $y = 0$, as shown in Fig. 3.3b. This shows the limit of no lensing when the source and the lensed misalignment is too large. The effect of the frequency is to increase the maximum value at $y = 0$ and increase the frequency of the oscillation. We can also see that the amplification is not infinite at $y = 0$, contrarily to what is observed in the geometrical optics approximation. In Fig. 3.3a, we can also notice that lensing has an effect up to a certain wavelength. When it is too large (so when $\omega \rightarrow 0$), we see that the amplification factor tends to unity.

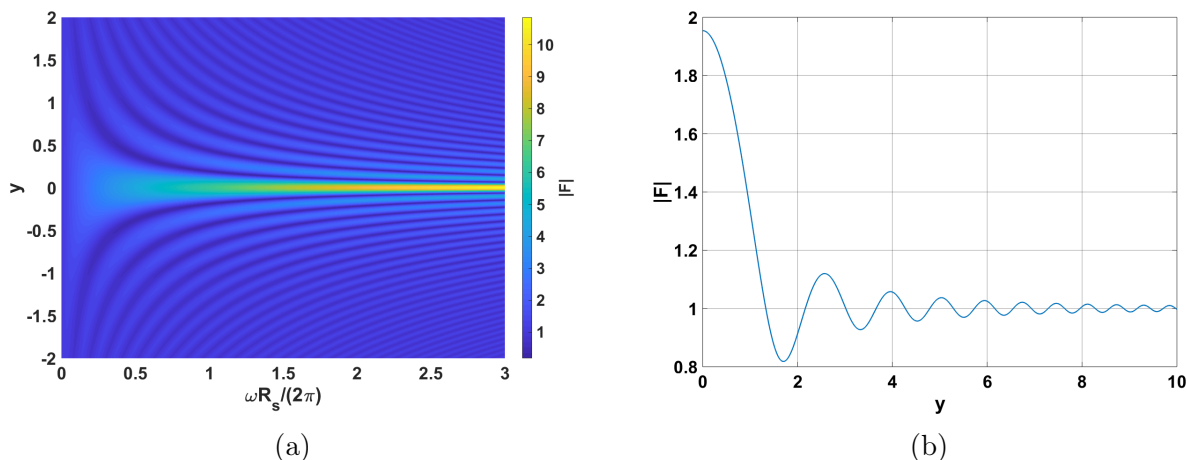


Figure 3.3: Evolution of the amplitude of the amplification factor F for $M = 60 M_\odot$ (a) as a function of y and ω (a similar figure can be found in [29]). (b) as a function of y with a fixed ω . Note that when $\omega R_s / (2\pi) = 1$, $\lambda = R_s$.

These figures, in particular Fig. 3.3b, show similar patterns as light diffraction and interference observed on a screen in slit experiments. We can observe spatial fringes as well, as

indicated by the variation of $|F|$ with y . It however oscillates around 1, because the initial wave is allowed to pass through the lens plane. The oscillations are then due to the interference between the two images at higher frequencies (geometrical optics approximation) and to diffraction at lower frequencies [29, 30, 38]. Indeed, let us recall that the integral form of $F(\omega)$ is (see Section 3.3)

$$F(\omega) = \frac{\tilde{\omega}}{2i\pi} V(\omega) = \frac{\tilde{\omega}}{2i\pi} \int_{\mathbb{R}^2} e^{i\omega \Delta t(\vec{x})} d\vec{x}. \quad (3.90)$$

At higher frequencies, we have seen in Section 3.3 that we can consider that there are multiple images, two for a point-mass lens (see Section 2.7). These then interfere to produce the observed fringes. At lower frequencies, the contributions from the whole lens plane matter, not only near the stationary points of $\Delta t(\vec{x})$. We can thus not say that ‘two images interfere’, it rather corresponds to diffraction.




Figure 3.4: Evolution of the amplitude of the amplification factor F for $M = 60 M_{\odot}$ as a function of ω (a similar figure can be found in [29]).

In Fig. 3.4, the evolution of $|F|$ is given as a function of ω for different values y . We see that at lower y , the maximum amplification is greater, as also expected from geometrical optics (see Section 2.7). The oscillations also have a larger amplitude and a lower frequency for small y . On the other hand, the amplification tends to 1 for any y at very low ω , as already mentioned. The oscillations also appear at lower frequencies for larger y . This suggests that geometrical optics is valid at lower ω for larger y . Indeed, these oscillations appear when the geometrical optics approximation of $|F(\omega)|$ becomes valid. This approximation with a point-mass lens, given in Eq. (3.32), reads

$$|F(\omega)| = \sqrt{|\mu_+| + |\mu_-| + 2|\mu_+|^{1/2} |\mu_-|^{1/2} \sin(\omega \Delta t_d)},$$

with Δt_d the time delay between the two images. The third term is thus the source of the oscillations that, hence, appear in the geometrical optics limit. These oscillations would then

be due to the interference between the two images, as already mentioned. The difference in the frequency of these oscillations can be explained by Δt_d . It is indeed lower for smaller values of y [38], hence the frequency of the oscillations of $|F(\omega)|$ is lower. At lower ω , diffraction occurs and the approximation Eq. (3.32) cannot be used anymore, which explains why the behaviour at low frequencies is different from that in the geometrical optics approximation. More complete discussions (and applications) can be found in [29], [30] and [38].

3.5.4 Validity and comments

First, we can show that this solution, and the equation in general, is valid only for point masses. The general gravitational potential reads

$$U(\vec{r}) = - \int \frac{G \rho(\vec{r}')}{|\vec{r}' - \vec{r}|} d\vec{r}', \quad (3.91)$$

so that in the approximation where $\vec{r}' \ll \vec{r}$, the potential reduces to

$$U(r) = - \frac{GM}{r}. \quad (3.92)$$

We thus retrieve the point-mass potential when considering the potential at a distance typically much greater than the size of the massive object, *i.e.* approximating the massive object by a point. This is in general not valid for galaxies. Let us consider the Milky Way, which has a mass of the order of $10^{12} M_\odot$, with M_\odot the solar mass. We then compute $R_s \approx 3 \cdot 10^{15} \text{ m} \approx 0.3 \text{ light-years (ly)}$. Let us assume that a similar galaxy acts as a lens on a wave coming from an infinitely distant source and that it is at a distance of 100 Mpc. We can then compute its linear Einstein radius $R_E^\infty = \theta_E^\infty D_L = \sqrt{2 R_s D_L} \approx 1.5 \cdot 10^4 \text{ ly}$, which is smaller than the radius of the galaxy. Since R_E represents a typical impact parameter, we cannot say that $r' \ll r$. This approximation is therefore more suitable for compact objects such as black holes. As a consequence, the solution that was derived is more appropriate to describe microlensing effects. For strong lensing by galaxies, $R_s \gg \lambda$, so the geometrical optics approximation should be appropriate to describe the phenomenon.

As mentioned earlier, the solution is valid for a plane wave, *i.e.* with a source at an infinite distance. A very similar solution can be derived with a spherical wave, as shown in [37]. In that case, we do consider D_S and D_{LS} in the expression of θ_E .

Another extension to the plane wave solution is to consider an expanding Universe. The metric to consider was already shown in Eq. (2.22). We can do a similar development as in Section 3.4 to derive the wave equation, with now

$$\sqrt{-g} = a^3 \sqrt{(1 + 2U)(1 - 2U)^3}, \quad (3.93)$$

and thus

$$\partial_\mu (\sqrt{-g} g^{\mu\nu} \partial_\nu \Phi) = \partial_\mu \left(a^3 \sqrt{(1 + 2U)(1 - 2U)^3} g^{\mu\nu} \partial_\nu \Phi \right) = 0. \quad (3.94)$$

Since a does not depend on the spatial coordinates, the corresponding terms will simply be multiplied by a . On the other hand, for $\mu = 0$, we also have the same term as before, but

multiplied by a^3 , and an additional term

$$3 a^2 \partial_0 a \partial_0 \Phi \sqrt{(1+2U)(1-2U)^3} = (1-2U)\sqrt{1-4U^2} 3a^2 \partial_0 a \partial_0 \Phi .$$

Then, similarly to (3.45) and dividing by a^3 , we get

$$-\frac{1-2U}{1+2U} \partial_0 \partial_0 \Phi - 3(1-2U) \frac{\partial_0 a}{a} \partial_0 \Phi - \frac{1}{a^2} \frac{4U}{1-4U^2} \partial_i U \partial_i \Phi + \frac{1}{a^2} \partial_i \partial_i \Phi = 0 , \quad (3.95)$$

where we recognise the Hubble parameter $H = \dot{a}/a$ from (1.23). The term containing the Hubble parameter can be neglected by noticing that the variation of a , of characteristic time $1/H_0$, is much lower than that of the wave. The other terms of order $O(U^2)$ can be neglected just as in the case of the non-expanding Universe. The equation then reads

$$\vec{\nabla}^2 \Phi - a^2(1-4U) \partial_0^2 \Phi = 0 . \quad (3.96)$$

We use the approximation that the derivative of a is negligible, so that we can set $\Phi(\vec{r}, t) = \phi(\vec{r}) \exp(-i\omega t/a)$. We then retrieve the same equation as the one we solved, but with $\omega \rightarrow \omega/a = \omega(1+z_L)$ (since the equation is solved in the reference frame of the lens). The solution is thus formally the same as the one that was derived (Eq. (3.88)) except that $\tilde{\omega}$ is now defined as

$$\tilde{\omega} = R_s \omega (1+z_L) . \quad (3.97)$$

In this case, one also interprets the distances as the angular diameter ones.

Finally, it is interesting to discuss how a given waveform is modified by lensing. If we define $\phi(\omega)$ as the frequency representation of the unlensed wave as seen by the observer (considered at a fixed position) and $\phi_L(\omega)$ for the lensed wave, we have, by definition of the amplification factor, $\phi_L(\omega) = F(\omega) \phi(\omega)$ [30]. The multiplication by $F(\omega)$ then produces particular patterns (such as oscillations) in the frequency representation of the lensed wave (see [29]). We can have a better representation of the modification of a waveform through lensing if we use the geometrical optics approximation of F (Eq. (3.30)). Indeed, in that context, we have

$$\phi_L(\omega) = \phi(\omega) \sum_j^N |\mu_j|^{1/2} e^{i\omega \Delta t_j^* - n_j i\pi/2} , \quad (3.98)$$

with Δt_j^* , n_j and μ_j respectively the time delay, the Morse index and the magnification of the j th image, and N the number of images. When we transform it back to the time domain, we find [30]

$$\Phi_L(t) = \sum_j^N |\mu_j|^{1/2} \Phi(t - \Delta t_j^*) e^{-n_j i\pi/2} . \quad (3.99)$$

This shows that, in the geometrical optics approximation and in the time domain, the amplitude of the wave, its phase and its arrival time are modified, but not necessarily the waveform. Indeed, if the differences in time delays are larger than the typical duration of a signal, then the images do not overlap and the lensed waveform is simply a succession of images of the initial waveform with different amplitudes (and phases). On the other hand, the superposition of the images (if the Δt_j^* are very close to each other) may significantly change the observed waveform.

Chapter 4

Gravitational lensing of gravitational waves: practical approach

Now that we have a good idea of the effect of the lens on the waveform, we can use our knowledge to explore how one may observe and identify two waves as being lensed. This is what will be done in this chapter. To be more precise, we shall focus on strong lensing, *i.e.* when multiple images of the same source are observed, with only a change in amplitude. I will first introduce the concepts required to understand the application, namely deep learning and gravitational-wave observation, detection and analysis. We will also discuss the advantages of deep learning applied to our problem, as well as a concrete example of a neural network that identifies lensed waves. Note that the notation for the vectors is changed in order to be more consistent with the usual notations used in machine learning and in the references.

4.1 Introduction to deep learning

Deep learning is one of the investigated techniques for identifying lensed events. It is an important concept in this thesis, which is why I will now introduce the basic concepts of supervised learning and deep learning. This section is largely based on [39], [40] and [41].

4.1.1 Supervised learning

Supervised learning is a statistical tool used to make inference. The idea is to use some data, called training data, that are representative of an unknown relation between some inputs and some outputs. These are used to estimate this relation, or *learn* it, and then apply this learned relation to make predictions on new inputs, as represented in Fig. 4.1. The goal is, of course, that the learned relation be as close as possible to the true relation. We shall see later how one measures the ‘closeness’ of the learned relation to the true one. Note that supervised learning differs from unsupervised learning, for which the true outputs are not known and are thus not part of the training data. The idea of supervision is that a ‘teacher’ provides the answer.

There are two main classes of problems that supervised learning can deal with: regression




Figure 4.1: General idea of supervised learning, taken from [41].

and classification. The former consists in predicting what a real output will be according to the inputs, *i.e.* output a quantitative prediction, *e.g.* fitting a function to a set of points. The second is concerned with attributing a class or label to the inputs, *i.e.* output a qualitative prediction. A typical example of this is to learn to recognise cats from dogs.

To be more formal, one can say that we consider that the training data (inputs \mathcal{X} and outputs \mathcal{Y}) are random variables distributed according to an unknown joint distribution $p_{\mathcal{X},\mathcal{Y}}$. In the case of regression, one tries to find the expected value $\mathbb{E}[Y|X = x]$, which is the best estimation for the output given the inputs [40]. As for classification, the attributed label is the one that maximises the probability of a label given the inputs. These expected values and probabilities are taken with respect to $P(\mathcal{Y}|\mathcal{X})$. The goal is thus to estimate this conditional distribution, learn it, so that when $\mathbf{x} \in \mathcal{X}$ is given, we may estimate the probability of the output $\mathbf{y} \in \mathcal{Y}$.

So, we want to find a relation f mapping \mathcal{X} to \mathcal{Y} , which is approximated by the function $\hat{f}(\mathbf{x})$ through the learning algorithm, where \mathbf{x} is a vector containing *features*, *i.e.* the inputs. All existing functions cannot be reproduced by a given algorithm. The set of functions that a learning algorithm is able to represent is called the *hypothesis space* \mathcal{F} . We consider that this algorithm has the ability to *learn*, *i.e.* to modify its approximation \hat{f} according to the discrepancy between the true output y_i , given in the training data, and its own prediction $\hat{f}(\mathbf{x}_i)$. We measure this difference using a *loss function* $L(y, \hat{f}(\mathbf{x}))$, which must always be positive (not strictly). It represents a penalisation on the error, so the algorithm will try to minimise it. We thus look for a function \hat{f} such that L is minimum, or to be more accurate, such that the expected loss is minimum. This expected value is called the *expected risk* $R(\hat{f})$, or generalisation error, which is minimised on the hypothesis space by the optimal function $\hat{f}_* \in \mathcal{F}$. We can thus write [39]

$$R(\hat{f}) = \mathbb{E}_{(\mathbf{x},y) \sim p_{\mathcal{X},\mathcal{Y}}} \left[L \left(y, \hat{f}(\mathbf{x}) \right) \right], \quad (4.1)$$

$$\hat{f}_* = \arg \min_{\hat{f} \in \mathcal{F}} R(\hat{f}), \quad (4.2)$$

where $\mathbb{E}_{(\mathbf{x},y) \sim p_{\mathcal{X},\mathcal{Y}}}$ denotes the expected value with respect to the distribution $p_{\mathcal{X},\mathcal{Y}}$. However, as stated, this distribution is unknown. We can estimate its value using an average over the whole training data set, since they are drawn from this distribution. This unbiased estimator

is called the *empirical risk* \hat{R} , or training error, and reads [39]

$$\hat{R}(\hat{f}, \mathbf{d}) = \frac{1}{N} \sum_{(\mathbf{x}_i, y_i) \in \mathbf{d}} L(y_i, \hat{f}(\mathbf{x}_i)), \quad (4.3)$$

where \mathbf{d} is the training data set containing all the input-output pairs (\mathbf{x}_i, y_i) . These are assumed to be independent and identically distributed. The optimal function over the data set may however not be same as the true optimal one, *e.g.* if the amount of data is not sufficient to allow an efficient training. The optimal function over the training set $\hat{f}_*^{\mathbf{d}}$ is such that [39]

$$\hat{f}_*^{\mathbf{d}} = \arg \min_{\hat{f} \in \mathcal{F}} \hat{R}(\hat{f}, \mathbf{d}), \quad (4.4)$$

It can be shown that $\hat{f}_*^{\mathbf{d}}$ converges towards the true optimal function \hat{f}_* as the number of data in the training set tends to infinity. Equation (4.4) represents the empirical risk minimisation principle.

There are a couple of interesting remarks to be made. First, we can see that the training data play a very important role, since the algorithm learns from them. The pre-processing as well as the very choice of the data is therefore critical and needs to be dealt with carefully. Then, we can notice that the set of functions that can be represented by a given algorithm is limited to its hypothesis space. It may therefore be that the optimal function \hat{f}_* , which minimises the risk among the functions of the hypothesis space, is not the optimal one when considering the set of all possible functions. As a consequence, it might be preferable to increase the number of functions that the algorithm can find to increase the likelihood of representing a ‘good’ function. This is done, for example, by increasing the number of parameters in the model.

However, increasing the capacity of the hypothesis space (roughly speaking, its size, or ‘ability to find a good model’ [39]) too much might not be a good idea. Indeed, when its capacity is too large, there can be functions that fit arbitrarily well the training data and one can thus face *over-fitting*. In that case, the model \hat{f} is so much adjusted to the training data that it is not able to correctly predict the output from the inputs that were not considered in the training procedure. It is important to keep in mind that we can only compute the *empirical risk* \hat{R} , but the true optimal functions are those that minimise R . In that case, the empirical risk is no longer a good estimator for the expected risk since it is biased towards the training data, so one cannot consider that the optimal function for \hat{R} is optimal for R . This is similar to fitting a noisy curve and having the solution pass through each and every point. This indeed minimises the error on the points, but not on the true unknown curve.

On the other hand, if the capacity of \mathcal{F} is too small, the algorithm can hardly represent a ‘good’ function, and the optimal ones in the hypothesis space lead to a relatively large expected risk. One talks about *under-fitting*. One looks for a balance between these two cases, such that the algorithm generalises well to data it was not trained on.

To see when a model is over-fitted or under-fitted, we need a test set, *i.e.* a data set containing data that are not in the training set. By doing so, we can evaluate how the algorithm performs with data it has never seen by estimating the expected risk on the data in the test set. Contrarily to the empirical risk in the over-fitting regime, this estimation is not biased, since \hat{f} was not adjusted to these data. It thus represents a better approximation

of the expected risk. It decreases in the under-fitting region, as better and better \hat{f} can be represented. It however increases in the over-fitting region since, as explained, \hat{f} does not generalise well. Therefore, the minimum of the generalisation error corresponds to an optimal capacity. Note that, in this process, the algorithm does not learn, *i.e.* it does not modify \hat{f} , since one evaluates its current approximation.

It has nevertheless been found recently that the generalisation error starts to drop again after a certain value of the capacity (in the over-fitting zone) and can even reach a minimum that is lower than the first one [42].

To end this introduction, it is important to talk about model selection and model assessment. As explained in [40], model selection is about testing the performance of different models in order to choose the best one, while model assessment is about estimating the generalisation error of the chosen model. We have already seen a criterion for choosing the best model, namely the one that minimises the generalisation error on a separate data set. This set on which one selects the optimal model is called the *validation set*. However, one cannot report the performance of the model on this set as its true performance, since it is biased, it was chosen according to its results on these very data. In a sense, it is then included in the training process, as it is used to make a choice on the model (*e.g.* on the architecture of the neural network). As a consequence, just as we need a validation set for training, we need a third data set, a kind of validation set for model assessment, called *test set*, in order to test the performance of the final model on data that have not been used in the former processes.

4.1.2 Deep learning

Deep learning is a class of methods of machine learning. As explained, machine learning allows one to learn transformation from inputs to an output. The raw data may not be suitable for learning such transformations, it may be better to use particular representations of the data for the algorithm to learn [41]. The idea of deep learning is to use several non-linear modules or layers that each after the other transform the representation of the data [43]. The representations themselves are learned by the algorithm, so that they finally amplify the relevant features and discard the variations that are unimportant to solve the task [43]. Neural networks implement such techniques and I shall only refer to them in the following. The rest of the section is mostly based on [41]. I thus recommend the interested reader to have a look at this reference to have more details and dive deeper into the subject.

Neural networks are then ‘mere’ non-linear statistical models. To start the explanation, it is easier to go for a linear regression problem. Let us consider a function $f : \mathbb{R}^n \mapsto \mathbb{R}$, and $y = f(\mathbf{x})$. We have measurements of N points with noise, such that the measured points are $y^{(i)} = f(\mathbf{x}^{(i)}) + \epsilon$. To approximate the function f , we search for a linear relation

$$\hat{y} = \mathbf{w}^T \mathbf{x} + b = \hat{f}(\mathbf{x}), \quad (4.5)$$

where \hat{y} is thus the prediction. In other words, we seek the weight vector \mathbf{w} and parameter b such that, in the framework of the least-square method, the sum of squared differences between $\hat{f}(\mathbf{x}^{(i)})$ and $f(\mathbf{x}^{(i)})$ is minimum, where the i index refers to the measured data. Note

that in this case, the sum of the squared differences is the loss function. The relation can in fact be represented by a single-layer neural network, as shown in Fig. 4.2a, where each line between x_i and \hat{y} is weighted. We thus interpret this graph as $\hat{y} = w_1 x_1 + w_2 x_2 + \dots + w_n x_n + b$, where b is implicit in the diagram and added to the output. The first layer is the input layer, since x_1, x_2, \dots, x_n are the inputs (components of \mathbf{x}), while the second one is the output layer. The input layer is not considered as being a layer of the network. This can easily be generalised to higher-dimension relations $\mathbf{f} : \mathbb{R}^n \mapsto \mathbb{R}^m$, with

$$\hat{\mathbf{y}} = \hat{\mathbf{f}}(\mathbf{x}) = \mathbf{W}\mathbf{x} + \mathbf{b} \quad (4.6)$$

$$\hat{y}_i = W_{ij} x_j + b_i, \quad (4.7)$$

where $\mathbf{W} \in \mathbb{R}^{m \times n}$ is the weight matrix and $\mathbf{b} \in \mathbb{R}^m$ is called the bias. The case where $m = 2$ is represented in Fig. 4.2b, but we can see that it can quickly become quite messy. As each input is linked to each output, it is called a fully-connected or dense layer. The goal of the network is to learn the value of the weights and the bias in order to minimise the loss (so the empirical risk).





Figure 4.2: Single-layer neural networks that represent the linear regression (a) with 1D output (b) with 2D output.

The next step is to show how we can get non-linear models. We can add to the networks of Fig. 4.2 a *hidden layer*, as shown in Fig. 4.3, *i.e.* a layer between the input and output layers. For obvious reasons, the weights are not represented on the diagram, but are still considered. One can even add several hidden layers, with p_i units each. This network architecture is called multilayer perceptron (MLP) and is the basic architecture.

Similarly to the single-layer network, we can easily interpret the diagram presented in Fig. 4.3. We can indeed first consider the hidden layer as the output of the input layer, and then as the input of the output layer. We still consider $\mathbf{x} \in \mathbb{R}^n$ and $\hat{\mathbf{y}} \in \mathbb{R}^m$ and introduce $\mathbf{h} \in \mathbb{R}^p$. We also need two different weight matrices, since there is one for going from \mathbf{x} to \mathbf{h} , and one from \mathbf{h} to $\hat{\mathbf{y}}$. We therefore introduce $\mathbf{W}^{(1)} \in \mathbb{R}^{p \times n}$ and $\mathbf{W}^{(2)} \in \mathbb{R}^{m \times p}$, and $\mathbf{b}^{(1)} \in \mathbb{R}^p$ and $\mathbf{b}^{(2)} \in \mathbb{R}^m$, and we have

$$\mathbf{h} = \mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)} \quad (4.8)$$

$$\hat{\mathbf{y}} = \mathbf{W}^{(2)}\mathbf{h} + \mathbf{b}^{(2)}. \quad (4.9)$$

Figure 4.3: Two-layer neural network, *i.e.* with one hidden layer.

If we replace \mathbf{h} by its expression in the second equation, we can directly see that the relation between $\hat{\mathbf{y}}$ and \mathbf{x} is still linear. However, the presence of a hidden layer leads to the introduction of non-linearities through an *activation function* σ , which is a non-linear function. Among them, we find the sigmoid, atanh or ReLU ($= \max(0, x)$) functions. The last one is generally preferred, since the two others are more prone to the vanishing gradient problem (see [41]). We define σ as the *element-wise* activation function, so that

$$\mathbf{h} = \sigma(\mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)}) . \quad (4.10)$$

We can inject this result in Eq. (4.9) to obtain a non-linear relation between $\hat{\mathbf{y}}$ and \mathbf{x} , which justifies the fact that neural networks are non-linear models. As already stated, we can still add several hidden layers, each with its own activation function, so that for q hidden layers, we have

$$\hat{\mathbf{y}} = \mathbf{W}^{(q)} \sigma_{q-1} \left(\mathbf{W}^{(q-1)} \sigma_{q-2} \left(\dots \sigma_1(\mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)}) \dots \right) + \mathbf{b}^{(q-1)} \right) + \mathbf{b}^{(q)} . \quad (4.11)$$

We can thus see that it gets very complex and highly non-linear. Neural networks therefore enable to represent complicated functions, but also simple ones. The price to pay is that they might require a lot of parameters (sometimes millions or billions nowadays), so that it is quite costly and requires specific hardware.

Finally, the last important point is the way that the network learns the different weights. As explained in the previous section, the network, through supervised learning, should be able to adapt its weights according to the loss, and the goal is to minimise the latter and, thus, the training error. Ultimately, though, it must also minimise the generalisation error. It is an optimisation problem, which is a bit more complex than usual ones, since it applies to a high-dimensional parameter space. It hardly admits an analytical solution and thus requires numerical methods. The first idea is to use gradient descent, *i.e.* we move in the

parameter space in the opposite direction of the gradient by steps proportional to its norm. We write this the following way [39]:

$$\theta_{t+1} = \theta_t - \gamma \nabla_{\theta} L(\theta_t), \quad (4.12)$$

$$L(\theta_t) = \frac{1}{N} \sum_{(\mathbf{x}_i, y_i) \in \mathbf{d}} L(y_i, \hat{f}_i(\mathbf{x}_i; \theta_t)) \quad (4.13)$$

where γ is the *learning rate*, θ_t is the value of the parameters at the step t and $\hat{f}_i(\mathbf{x}_i; \theta_t)$ stresses the fact the \hat{f} is defined by the parameters θ_t . In our case, θ contains the value of all the weights. However, given the large number of parameters, it is in practice impossible, or much too costly, to compute the gradient with respect to all parameters. Other optimisation algorithms that approximate this method and are faster to apply are used in practice. The idea remains nevertheless the same, and *automatic differentiation* is used to efficiently compute the gradient (or its approximation). I will not detail this technique here, but more information may be found in [44]. Once the gradient is computed, one may move in the parameter space according to (4.12), which means that the weights are updated. This is how the network can learn and change its approximation \hat{f} .

The value of the learning rate is also critical to ensure convergence. If it is too small, one may end up in a local minimum without being able to escape it, while if it is too large, one may never find a minimum or oscillate around it without ever reaching it.

4.1.3 Convolutional Neural Networks (CNN)

The MLP is one architecture among many different ones. Convolutional Neural Networks (CNN) are another popular architecture for networks. These are typically used on images, or time series, to extract *features*, or structures in the data. This is thus the usual network used on image classification problems, since we expect each class (*e.g.* cats and dogs) to have proper features that may allow one to distinguish them. Similarly, we know that gravitational waves signal display a specific pattern of increasing amplitude and frequency, so CNNs may well be useful to detect them. It can also be noted that CNNs must be (and are) invariant under translations, *i.e.* it must respond the same way to the same feature of the input regardless of where it is within it. In other words and for example, it must be able to detect a gravitational wave whether the merger be at the beginning of the time series or at the end. The idea is therefore to look at the input only locally. This way, if a local pattern is moved within the input, it does not matter where it is located, since the network will only treat the local pattern itself and not the pattern and its environment. The discrete convolution implements this idea of locality.

In fact, CNNs use cross-correlation and not convolution. Let us consider the 1D case. Even though it can easily be generalised to higher dimensions, we will not need it later. Let us introduce $\mathbf{x} \in \mathbb{R}^W$ and $\mathbf{u} \in \mathbb{R}^w$, with $w \leq W$. The discrete cross-correlation of \mathbf{x} with \mathbf{u} is written as [39]

$$(\mathbf{x} \star \mathbf{u})[i] = \sum_{j=0}^{w-1} x_{i+j} u_j. \quad (4.14)$$

Therefore, the resulting vector has a size $W - w + 1$. This operation can be viewed as sliding the vector \mathbf{u} across \mathbf{x} by step of 1, and at each step computing the element-wise product of \mathbf{u} with the corresponding w components of \mathbf{x} it overlaps. For example,

$$\mathbf{x} (1 \ 0 \ 3 \ -1 \ 1 \ 2) \quad (4.15)$$

$$\mathbf{u} (1 \ 2 \ 1 \ 0) \quad (4.16)$$

$$\Rightarrow (\mathbf{x} \star \mathbf{u}) = (1 \cdot 1 + 2 \cdot 0 + 1 \cdot 3 + 0 \cdot (-1) \quad 1 \cdot 0 + 2 \cdot 3 + 1 \cdot (-1) + 0 \cdot 1 \quad \dots) \quad (4.17)$$

$$\Rightarrow (\mathbf{x} \star \mathbf{u}) = (4 \ 5 \ 2) . \quad (4.18)$$

We thus see that u provides weights to perform an operation, or linear transformation, on a part of the components of \mathbf{x} . These coefficients do however not change, so the transformation is the same everywhere! We can see how this relates to the idea of locality, it does not matter where a feature is in \mathbf{x} , it will be found whatever its location, since the same transformation will be applied to it regardless of its where it is.

In general, \mathbf{x} is called the input feature map, \mathbf{u} the *kernel* and their cross-correlation is the *output feature map*. Also, w is the receptive field. It can be shown that this architecture is in fact a special case of a fully-connected layer. Intuitively, the weights will always be the same, corresponding to the value in the kernel, and will be zeros outside the kernel size. This is illustrated in Fig. 4.4, where $\hat{\mathbf{y}} = \mathbf{x} \star \mathbf{u}$ and u_i is the value of the i th component of kernel, so a weight. The idea of locality is even more obvious in that diagram.




Figure 4.4: Illustration of a CNN as a fully-connected layer. The weights are null when there is no link.

One can also take into account the *channels*, e.g. the RGB ones for images. In that case, we have $\mathbf{x} \in \mathbb{R}^{C \times W}$ and $\mathbf{u} \in \mathbb{R}^{C \times w}$, where C is the number of channels. contrarily to a multidimensional convolution, we usually do not consider that the kernel ‘slides’ along the

channel dimension. As a result, we simply sum over the values of the channels, so that the output will have 1 channel (so it is a vector). Mathematically, it reads [39]

$$(\mathbf{x} \star \mathbf{u})[i] = \sum_{j=0}^{w-1} \sum_{c=0}^{C-1} x_{c,i+j} u_{c,j}. \quad (4.19)$$

We can however still obtain multiple channels as output by performing the ‘convolution’ several times with different kernels.

When implementing the CNNs, one can specify the padding. The latter adds 0 elements at both ends of the vector (so a padding of n adds $2n$ null elements). This allows, among other things, to conserve the size of the input throughout the convolutions.

The last point to discuss is *pooling*. This reduces the size of the input of the convolution by averaging over or taking the maximum among p elements. The number of elements is thus, roughly, divided by p . Though we might lose a bit of information, the input keeps its global structure. The advantages are, of course, that the number of parameters decreases and the computation cost can be greatly reduced, but also the size of the effective receptive field increases. The effective receptive field of a layer is all the elements of previous layers that are involved in the computation of its output [41]. For example, if we apply a pooling of size 3 on a signal, then a convolution with a kernel size 3, the receptive field on the convolution input is 3, but 9 on the initial input, since elements were gathered by groups of three. Therefore, the first layers will search for features at small scales, while the last ones will rather identify features in a broader area. It is nonetheless worth mentioning that CNNs do not always see features as we do, or would want them to. It may not be easy to interpret what the features identified by the CNNs are.

A CNN is generally defined as a composition of convolution, activation function, pooling and fully-connected layers. Typically, convolution is applied, followed by an activation function. The result is then pooled, and this can be repeated several times before the output of this part is fed to one or several fully-connected layers.

4.2 Gravitational-wave observations

Before diving into the application of neural networks to gravitational-wave data analysis, one must first introduce how these waves are observed, what a typical wave signal looks like, and what the main challenges in detecting a gravitational wave are.

First, it must be noted that the only type of gravitational wave that we have been able to detect so far are those coming from compact binary coalescence (CBC), *i.e.* mergers of binary black holes, neutron stars, or a neutron star with a black hole. The signal is characterised by an increase in frequency with time. Intuitively, the radiation is emitted as the bodies circle around each other. This emission results in a loss of energy, transported by the wave. As a consequence, the two objects get closer to each other, which increases the frequency of the orbit and hence of the emitted wave, until they merge. An example is shown in Fig. 4.5. The strain is what is measured by the detectors, which are laser Michelson-like




Figure 4.5: Example of a gravitational wave from a merger of two black holes of equal mass $M = 50 M_{\odot}$ at a distance of 100 Mpc. Computed with PyCBC [45] with the SEOBNRv4_opt model [46, 47].

interferometers. The current ones, LIGO, Virgo, KAGRA and GEO 600, are L-shaped with arms of a length of 3 to 4 km [5–9] (600m for GEO 600 [10]). When the gravitational wave passes through the detectors, the mirrors move, which introduces a change in the interference pattern. The strain corresponds to the ratio of the change in the length of the arms (due to the motion of the mirror) over the total arm length. In the example of Fig. 4.5, the strain of 10^{-20} is relatively large, the typical order of magnitude is 10^{-21} . In fact, if we look at the problem in the TT gauge (see Section 3.1.3), the mirrors do not move, *i.e.* they keep the same coordinates, but the distance travelled by light is still different from that without the gravitational wave, since the metric changes according to Eq. (3.18) [27]. The results are the same, only the interpretation changes. It is also important to note that the sensitivity of the detectors is not the same in all directions.

An important characteristic of the observation data is the noise. The measured quantities are indeed so small that any small perturbation may cause a significant noise level, even quantum effects. There are several sources of noise: the temperature of the mirror, the irregular arrival of photons on the mirror (shot noise), radiation pressure that slightly moves the mirror, the power fluctuation of the laser, the seismic activity of the Earth (and human activity), and a few others [6, 48]. The noise does not have the same amplitude at all frequencies, it is not white. The typical power spectral density (PSD) is shown in Fig. 4.6. This reduces the possibility of detection in the range 20-1000 Hz, otherwise the noise amplitude is too large compared to the signal. In data processing, it is common to apply a band-pass filter to discard the parts of the strain that are largely dominated by the noise, and to whiten the noise. This is done by dividing the Fourier spectrum of the data by the square root of the PSD. The resulting signal is then characterised by a white noise and the wave signal should have a larger amplitude relative to the noise. It remains nonetheless that the noise is generally much larger than the signal. This is the main challenge in detecting




Figure 4.6: PSD of aLIGO detector for the O3 run. Computed with PyCBC.

gravitational waves. The increase in the sensitivity of the detector, hence the reduction of the noise, allows one to detect more events, and weaker ones.

It is important to note that the noise is neither Gaussian nor stationary. It is however often assumed to be both when starting to try some networks to analyse the data. In true noise, there are what one calls glitches, short bursts of noise that can have different characteristics.

4.3 Importance of artificial intelligence in gravitational wave detection

I will now expose some reasons why deep learning may be an important tool in gravitational-wave data analysis. To this end, one first needs to introduce the current main method along with its advantages and flaws. Then, we will see how deep learning may help and I will briefly present what is currently done in that field.

4.3.1 Detection of gravitational waves: matched filtering

The most common technique used to detect gravitational waves is matched filtering. Recently, artificial intelligence has also been investigated for this task. Matched filtering requires to use templates of signal to search for a similar waveform within the observation data. It can thus only be used to detect modelled waves and takes some time to give results, since one needs to try many different templates on long signals. I will not develop much the theory of matched filtering, but more information can be found in [48–50].

The idea of matched filtering is to compare the data with a template using cross-correlation. The computed cross-correlation would peak where the data and the template match. Let us consider the output of a detector, *i.e.* strain data, $s(t) = n(t) + h(t)$, with $n(t)$ the noise (assumed to be Gaussian and stationary) and $h(t)$ the gravitational-wave signal.

We also consider a template $u(t)$. One computes the SNR, defined as [50]

$$\frac{S}{N} \equiv \rho \equiv \frac{\langle s, u \rangle}{\text{rms} \langle n, u \rangle}, \quad (4.20)$$

where rms is the root mean square error and the inner product is the noise-weighted cross-correlation

$$\langle s, u \rangle = 4 \text{Re} \left[\int_0^\infty \frac{S^*(f) U(f)}{S_n(f)} df \right], \quad (4.21)$$

with S and U the Fourier transforms of s and u , respectively, and S_n the one-sided power spectral density of the noise. It can be shown that the denominator is equal to $\sqrt{\langle u, u \rangle}$ [49, 50]. This SNR is maximum when the template matches the strain data, which suggests that a wave is detected. It is indeed unlikely to have a very large SNR if $s(t) = n(t)$. One can then define the optimal SNR (SNR_{opt} or ρ_{opt}), which is obtained for $h(t) = Au(t)$, such that [51]

$$\mathbb{E} [\rho_{\text{opt}}] = \mathbb{E} [\langle n + Au | u \rangle] = A, \quad (4.22)$$

with A the amplitude and \mathbb{E} the expected value (over an infinite number of noise realisations). This is true provided the noise has a null expected value, unit variance and that $u(t)$ is normalised such that $\langle u | u \rangle = 1$. If the latter condition is not fulfilled, one can still define [52]

$$\sigma^2 = \langle u | u \rangle = 4 \int_0^\infty \frac{|U(f)|^2}{S_n(f)} df, \quad (4.23)$$

and divide u by this quantity to normalise it.

A significant SNR is however not always sufficient, since the noise amplitude is in general of the order or greater than the wave amplitude. In practice, the strain data of all the available detectors are used to make the detection more significant. If there is indeed a wave, one can estimate the parameters of the source according to the waveform. The parameters influencing the waveform with CBC waves are the masses of the compact bodies, their spin vector and orbital angular momentum, the luminosity distance (1.34), the position in the sky (right ascension, declination), as well as the time and phase of merger [48]. To be more accurate, these are the parameters for a binary black hole merger with a quasi-circular orbit. There are additional parameters when one considers eccentric orbits and neutron stars. These parameters can be entered as arguments of template generation functions such as those provided by the python software `PyCBC` [45], which uses generators/*approximants* implemented by the `LALSuite` software [46].

The generation of such theoretical waveforms demands (sometimes heavy) numerical computations from the equations of general relativity. Besides the cost of building the templates, the matched filtering requires a certain number of them in order to span (part of) the parameter space and, thus, make sure that if a wave is present in the observation data, it will produce a significant SNR and will be detected. The true parameters of the wave are indeed unknown, so we have to try many different templates to find the most suitable one. It is however not necessary to have one for each and every possible combination of the parameters (which would be infeasible), since some lead to very similar waveforms. Intuitively, we can say that the number of templates needed not to miss a wave is large,

given the number of parameters. The number of operations to test N templates on a data segment is proportional to N (and also depends on the length of both the templates and the data) [50]. This, added to the time to generate templates, gives an idea on the possibly large cost of this technique. Less accurate templates can nonetheless be used for a faster ‘pre-identification’, before a confirmation [53]. It seems then natural to say that this method is costly and, therefore, time-consuming. This is even more apparent when we know that we have more and more data to analyse. As this happens and the sensitivity of the detectors increases, the matched filtering method may become too slow to detect and analyse each and every possible wave hidden in the data. It remains nevertheless that matched filtering is the optimal method in stationary, Gaussian noise [48, 50, 53].

4.3.2 Deep learning: filtering and early warning

Given the cost of matched filtering and the increasing number of detectable events, the need for faster algorithms becomes a serious issue. In fact, one does not especially need more accurate algorithms, as long as they are less costly. There are two main reasons for searching for faster methods.

The first one is to find, with such methods, potential events in the observation data. These candidates can then be confirmed or discarded by more accurate and time-consuming methods. This prevents wasting time trying to identify signals where there are probably none. The second reason is to make early-warning detections [54]. It consists in detecting a coalescence signal fast enough to warn observatories that there might be an electromagnetic counterpart of the gravitational wave to observe. This is of interest for mergers which involve a neutron star. The interest in observing an electromagnetic counterpart has already been demonstrated. One of the neutron star mergers that was detected, GW170817, was found to be almost coincident with a gamma-ray burst, which arrived some 1.7s after the gravitational-wave signal [55]. The joint observations were used to set constraints on the speed of gravity, very close to c , which rules out some alternative theories of gravity. It was also used to test general relativity and study the gamma-ray burst. Finally, it allowed one to estimate the Hubble constant, from the redshift measurement of the host galaxy and the estimated distance from the gravitational-wave signal [56]. The value is however not accurate and one would need more such joint observations to obtain a more accurate estimation of H_0 . To sum up, such multimessenger detections have a real scientific interest, which stresses why early warning is important.

In the case of lensing, one would mainly benefit from the filtering advantage, which is, in fact, even more important in this case. Indeed, the number of pairs to test increases roughly as the square of the number of detected events. Since the latter and the rate of detection will increase (because of new generation and/or improved detectors), it will probably be infeasible to keep up and analyse all the pairs with the current costly methods, which will be presented in the next section. This becomes even more critical if one wants to find triplets or four images of the same source, since the number of groups of events to analyse is even larger.

Deep learning may provide us with such faster algorithms. Indeed, even if the training takes some time, the predictions of the network are in general very fast. Also, it was shown that matched filtering is equivalent to a specific neural network [57]. This may be no proof

that it is the best solution, but it is still, in my opinion, a clue that machine learning is an interesting and promising lead to follow.

Among the works that apply machine learning to the detection of gravitational waves, some analyse directly the time series with a 1D CNN, achieving similar performance to matched filtering in Gaussian noise [58], but also in true LIGO noise [59]. Some networks also estimate the value of the parameters on top of detecting the wave [59] and, in addition to this, some also classify the event as binary black hole or neutron star coalescence [53]. Another way to look for gravitational waves is to analyse the spectrogram of the signal [60].

These results are encouraging. Some are indeed obtained under idealised conditions, but some also under realistic ones such as true noise with glitches. There are also other investigated applications of machine learning in gravitational-wave data analysis. We can mention the extension of the detection to unmodelled waves [61], the identification and classification of glitches, subtraction of noise, and waveform modelling [16]. These are not really relevant to discuss in this thesis, but the interested reader may find a review in [16].

4.4 Identification of lensed gravitational waves

Bayesian inference has been used to detect lensed events [12], and artificial intelligence also starts being investigated [62, 63]. We shall briefly see what are the basic principles of these techniques and present the current status of the search for lensed signals.

We shall first consider strong lensing by galaxies or galaxy clusters. Bayesian methods are based on a posterior overlap of the parameters [64]. The idea is that, under the lensing hypothesis, the overlap between the parameter posteriors of both waves should be large, since strong lensing only modifies the amplitude and the phase. This method can be used to keep only the most promising pairs, as was done in the search for lensing events in [12]. The verification of these pairs is then made with a joint parameter estimation analysis. This allows one to compute three statistics to evaluate the likelihood of the lensing hypothesis. The first one takes into account the proximity of the parameters of the two waves. The second one is the same as the first, except that it takes into account the population of sources and lenses, in order to account for the fact that two sources with similar characteristics and close in the sky may emit a similar wave that can then look like a pair of lensed waves, though it is not. A third quantity takes into account the fact that some waves are preferentially detected in some mass ranges and sky localisation because of the sensitivity of the detectors [12]. More details can be found in [12] and references therein.

As for deep learning methods, [62] uses a DENSENET, which is a CNN initially designed for image recognition [65]. This network is used on the Q-transform (frequency vs time) of the waves. The idea is thus to compare the frequency evolution of both waves to see if they are identical or not. The sky localisation is also used in the process and the detectors are treated separately, each corresponding to one network. The data are generated according to a given population of sources and a given lens model. The noise that is used is Gaussian and stationary. This model is more complex than what will be considered in the next section. Also, the representation of the data will not be the same.

Similarly to strong lensing by galaxies, we can search for microlensing signatures including wave effects as discussed in Chapter 3. In this case, we generally do not have multiple images but a deformation of the waveform due to wave effects or superposition of multiple images with a time delay smaller than the duration of the wave signal (see [12] and references therein). Deep learning was also tried to identify such events. For example, [63] uses a CNN on the spectrogram of waves to identify lensing signatures, such as beating patterns.

The LIGO Scientific Collaboration and the Virgo Collaboration have searched for lensed events among the detected ones in [12]. This paper only reports on the analysis for the O3a run, and one still has to wait for the results of O3b. They conclude that no event can be considered to be lensed or multiply-imaged with enough confidence. As a result, no positive detection of lensing has been made yet.

4.5 A neural network to find them

In this section, I shall propose a neural network model to identify lensed waves. The considered problem will first be properly described, and then the data generation and network architecture will be presented. Finally, the results will be discussed. Potential improvements to the model will be envisaged in the last chapter.

4.5.1 Formulation of the problem

The model aims at distinguishing pairs of events that are lensed from those that are not. We thus consider strong lensing, *i.e.* with multiple images of the source. The sources are exclusively binary black holes. As already mentioned, geometrical optics is sufficient to describe strong lensing by galaxies or clusters of galaxies, so that one can consider that the only difference between two images is their amplitude and phase. We shall however consider that the phase remains the same. We shall also use a Gaussian stationary noise and use only a single detector. Contrarily to previous works, the network analyses directly the time series. These are nevertheless modified by different processes described in the next section. To my knowledge, no other work using the time series for this problem has been published. This model is a mere proof of concept, to show that using such a representation of the data can accomplish the task. More work will be needed to test the performance of such networks further using more complex data. We shall focus on a rather simple case where no model is used for the lens and where we do not take into account any time delays, accurate magnifications, lensing probabilities or parameter distribution of the source population. We shall consider that one wave is magnified and the other demagnified.

The goal is, of course, to identify as many lensed events as possible. There is an additional constraint that the false positive probability (FPP, ratio of wrongly classified unlensed events to the total number of unlensed events) be as low as possible. Indeed, if the network is to be used to identify candidates that will be checked by time-consuming methods, one would like to have a minimum of false positives not to lose a lot of time with the other methods. This aspect is important to keep in mind for testing the network. Of course, a lower FPP also increases the confidence in the predictions.

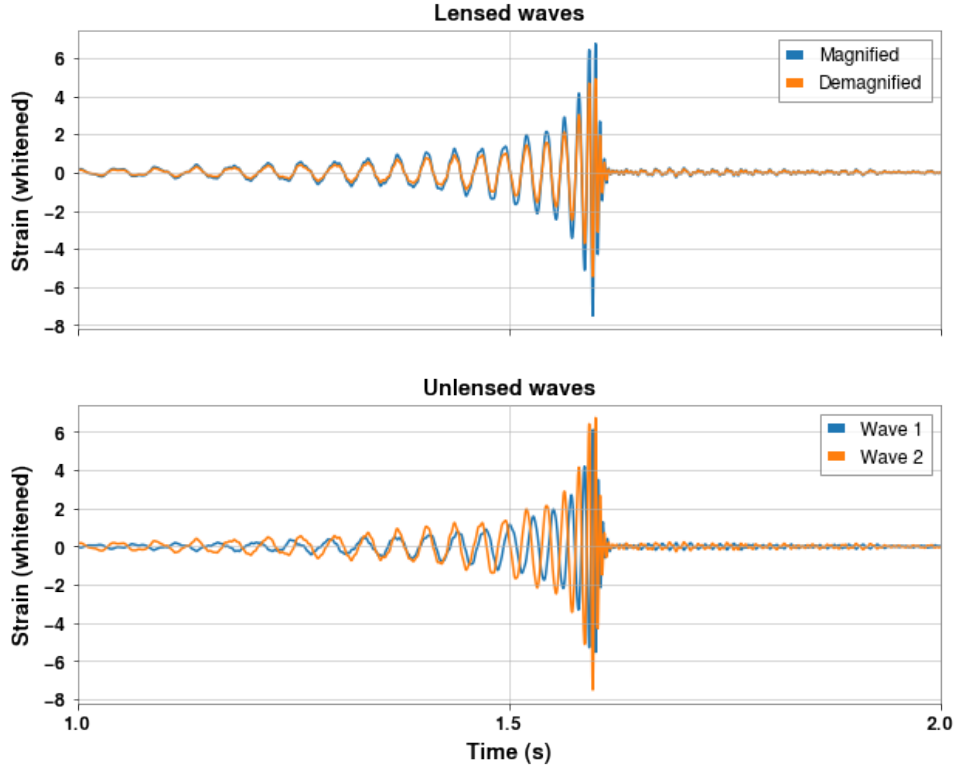


Figure 4.7: Example of lensed and unlensed (processed) waves superimposed. Data generated and processed with PyCBC, using the *SEOBNRv4_opt* model for templates.

It may be interesting to have a look at what the neural network must distinguish. This is depicted in Fig. 4.7. One can see that, as stated, only the amplitude changes between the two images in the lensed case, while the two images differ in both amplitude and frequency evolution in the unlensed case. This may look rather simple to distinguish, but once the waves are injected into the noise, the task can quickly become difficult, in particular for low SNRs. Also, the two waves in the unlensed case can be very different one from the other or have very similar parameters, hence waveforms, which makes it difficult to tell them apart. Note that these waves are not fed to the network as such, since they still need to be added to noise with a given SNR.

4.5.2 Data generation

Data were generated using dedicated methods in PyCBC [45], using the *SEOBNRv4_opt* [46, 47] approximant. The spin and inclination of the black holes are neglected. The only relevant parameters to enter are the masses of the two merging black holes, the location of the source in the sky and the polarisation angle. The distance to the source is not important, since the wave will be re-scaled. I considered a range of chirp masses (see Eq. (5.1)) between 15 and 40 M_{\odot} and a mass ratio m_1/m_2 between 1 and 4. These parameters correspond to a total mass of the system between around 30 and 120 M_{\odot} . Both are drawn uniformly in their respective interval. The range of chirp masses is not based on an accurate population

estimate, but roughly estimated from detections of the O3 observing run [4, 66], for it is more likely to find a pair of events in that run than in previous ones (given the larger number of detections). The minimum is however larger in order to have a limited duration of the generated event. The mass ratio range was chosen arbitrarily. As for the position of the source and the polarisation, they were set to a constant which is the same for all generated waveforms. These were chosen, along with the projection time onto the detector, such that the projected wave does not have a null amplitude. It is important to note that unlensed events are generated one at a time, while lensed events are generated in pairs. Each event of the pair is however treated separately, each is associated with a different noise realisation.

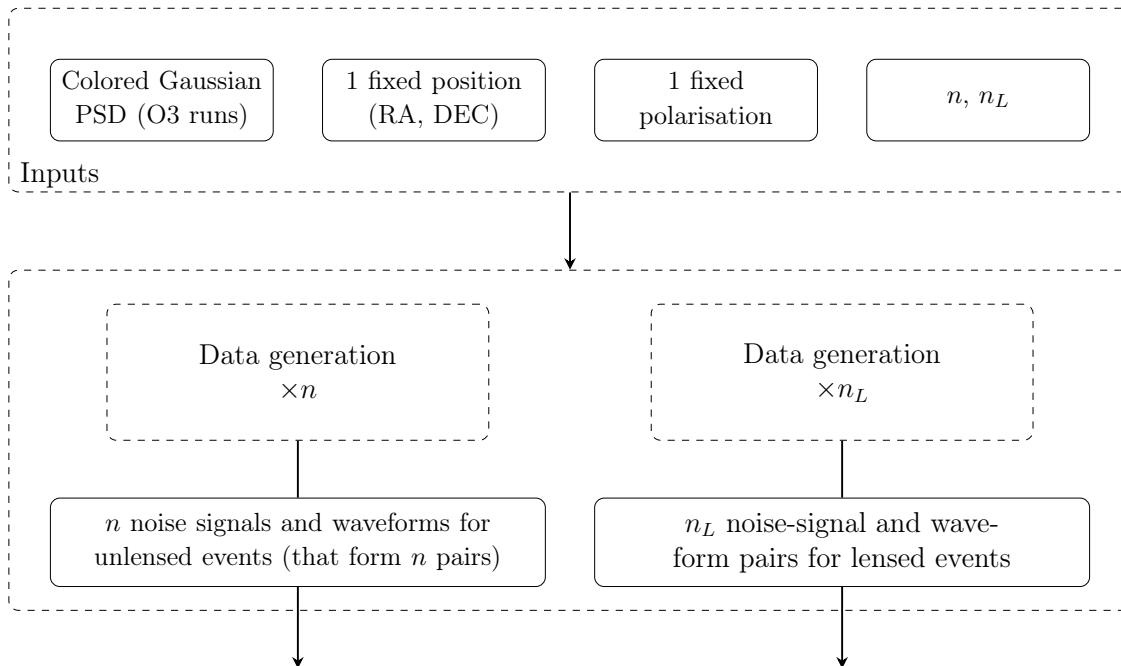


Figure 4.8: Diagram representing how the data set is generated.

As for the noise, it was generated according to the PSD of the noise for the O3 run, more precisely with a BNS range of 140 Mpc [46, 67] (range at which the merger of two neutron stars of mass $1.4 M_{\odot}$ can be detected with an $\text{SNR} \geq 8$ [5]). The obtained noise is Gaussian and stationary, contrarily to the true noise. This is however sufficient for the purpose of this work.

Once the waveform and the noise are generated, they both need to be processed. This processing is summarised in Fig. 4.8 and Fig. 4.9. The position of the source, polarisation of the wave and PSD of the noise are the same for all the events that are generated. They are used to produce n pairs of unlensed events and n_L pairs of lensed events, as shown in Fig. 4.8.

In Fig. 4.9, the data generation process is summarised. First, the wave is projected onto a detector (arbitrarily LIGO Hanford) at a fixed detection time t_0 corresponding to the beginning of O3. Afterwards, the merger time, associated with the peak amplitude of the wave, is placed at a time $t = 1.6$ s, so that enough of the pre-merger and post-merger phases can be represented. Aligning at merger time is important for the network not to rely on

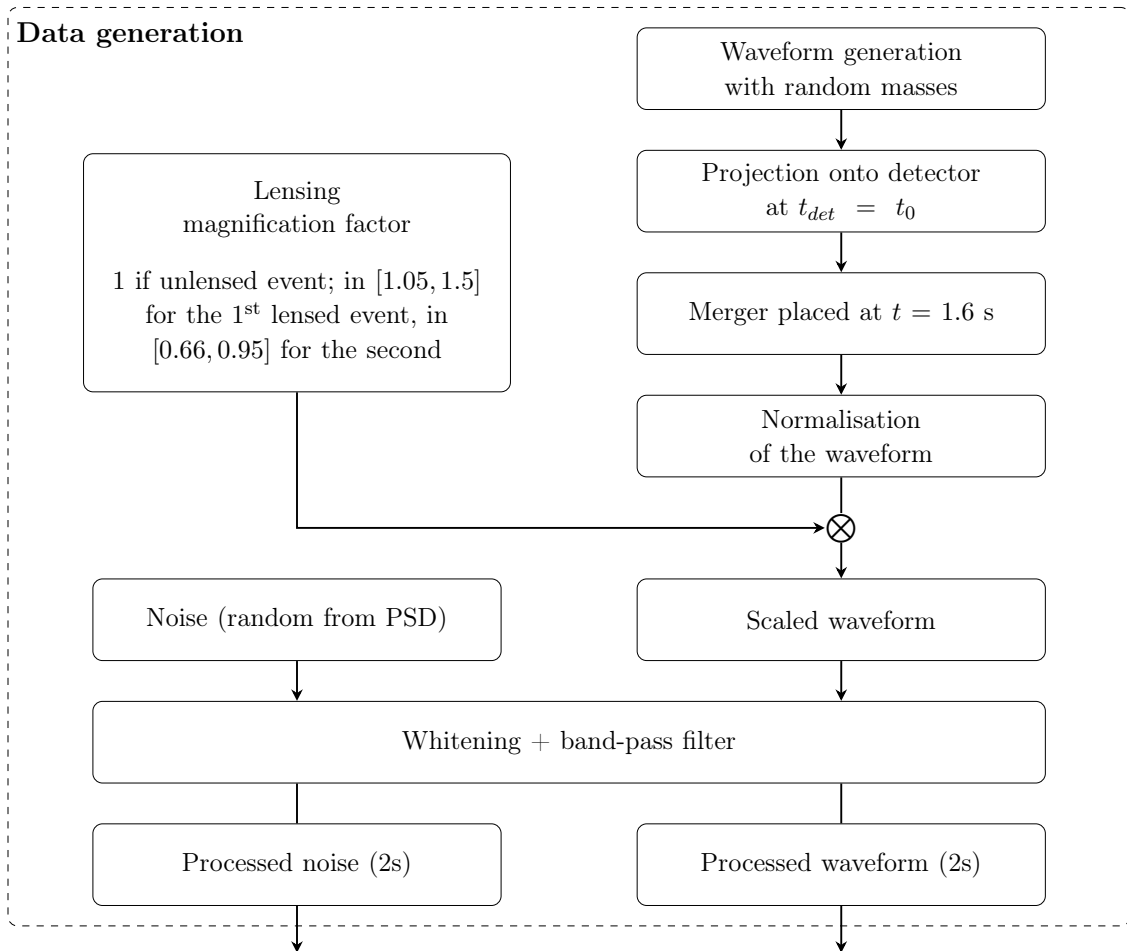


Figure 4.9: Diagram representing how the noise signals and waveforms are generated and processed.

features arising from a misalignment, which are irrelevant as they are independent of the wave itself. After that, the waveform $u(t)$ is normalised such that $\langle u|u \rangle = 1$. This is done by dividing the waveform by σ as defined in Eq. (4.23). This allows us to get the optimal SNR, since the waveform is also the template and will be injected into the noise, so the gravitational-wave signal and the template are the same (see Section 4.3.1 and Eq. (4.22)). This quantity will be used to measure the importance of the noise with respect to the wave. In the following, the SNR refers to the optimal one. Then, the wave amplitude is further scaled by a lensing magnification factor. If one produces an unlensed event, there is no magnification, so the factor is 1. In the case of lensing, events are generated in pairs, one of the events being magnified with a factor randomly drawn in $[1.05, 1.5]$ and the other being demagnified with a factor randomly chosen in $[0.66, 0.95]$. Note that this scaling modifies the optimal SNR, which is then multiplied by the lensing magnification factor. This allows one to have different SNRs for both images, which is more realistic. However, if one takes into account the rotation of the Earth and the time delays, the difference in SNR can be lower or larger than that induced by a difference in magnification because the orientation

of the detector will have changed between the two detections. The ranges of factors were chosen such that the difference in SNR is not too large and for the greatly demagnified wave to still have a reasonable SNR.

After all these steps, both the noise and the waveform are whitened by dividing their Fourier transform by the square root of the PSD. A band-pass filter is also applied to keep the frequencies in $[20, 2000]$ Hz. They are then cropped to a duration of 2s, at a sampling rate of 4096 Hz.




Figure 4.10: Example of a whitened waveform in orange and the corresponding signal (noise + wave) in blue seen at different optimal SNRs.

An example of data is shown in Fig. 4.10 where the waveform (already processed) is superimposed to the corresponding signal (noise + wave) at different SNRs. We can see that the wave can hardly be seen in the signal below an SNR of 20, even though we can guess its presence at an SNR of 10 because we know where the wave is and what it looks like.

The noise and waveforms are added together when loading the data set. The advantage of doing so is that we can multiply the amplitude of the wave by a given factor in order to obtain the desired optimal SNR without having to generate a data set for each of these SNRs. Indeed, given Eq. (4.22), we get the optimal SNR A if we multiply the waveform amplitude by A and inject it into the noise, since the waveforms we use are normalised.

4.5.3 Model and training

The architecture is presented in Tab. 4.1. I used a CNN with an increasing number of output channels, a kernel size of 3 and a maximum pooling with a size of 4. The relative

simplicity of the models proposed by [58] and [59] for gravitational wave detection oriented the choice towards simpler models. Preliminary tests suggest that deeper and more complex networks do not yield especially better results. This should however be verified in a more formal manner.

	Layers	Kernel size	Input size	Output size
0	1D convolution	3 (padding 1)	2×8192	64×8192
1	ReLU	/	64×8192	64×8192
2	Max pooling	4	64×8192	64×2048
3	1D convolution	3 (padding 1)	64×2048	128×2048
4	ReLU	/	128×2048	128×2048
5	Max pooling	4	128×2048	128×512
6	1D convolution	3 (padding 1)	128×512	256×512
7	ReLU	/	256×512	256×512
8	Max pooling	4	256×512	256×128
9	1D convolution	3 (padding 1)	256×128	512×128
10	ReLU	/	512×128	512×128
11	Max pooling	4	512×128	512×32
12	1D convolution	3 (padding 1)	512×32	1024×32
13	ReLU	/	1024×32	1024×32
14	Max pooling	4	1024×32	1024×8
15	Flatten	/	1024×8	8192
16	Linear	/	8192	50
17	ReLU	/	50	50
18	Linear	/	50	25
19	ReLU	/	25	25
20	Linear	/	25	1
21	Sigmoid	/	1	1

Table 4.1: Architecture of the neural network with the kernel size and output channels for the convolution part, and the input and output units for the MLP .

The training was done on 10^5 pairs, with an equal number of lensed and unlensed ones. The loss was chosen to be the binary cross-entropy and the optimiser is Adam [68] with its default parameters. Both training and testing were performed on an NVIDIA GeForce RTX 2070 SUPER GPU. At each epoch, the network is tested on a validation set containing 5×10^4 pairs in the same proportion as in the training set. The best one is kept for testing on the test set containing the same amount of data as the validation set. The latter was also used to choose the best model among the ones that were tested. Splitting the data into those 3 sets should be sufficient, for the amount of data is large enough to span uniformly the whole parameter space in the considered ranges and in each of the sets. Also, the learning rate starts at 6×10^{-5} and is divided by 2 when the validation loss decreases by less than 5% or increases. It stops decreasing once it is smaller than 6×10^{-6} . The batch size is 500.

Also, the true number of data on which the network is trained is twice as large as the size of the training set. The reason is that each pair is used twice, once with the magnified event first and then the demagnified one, then the same pair but reversing the order of the events. The purpose is to ensure that both channels are trained the same way. Since they have

no physical nor relevant meaning in the context of this work, there should be no difference between the two, *i.e.* the results should be the same whether an event is placed in one channel or in the other. This might not be the case when considering time delays, provided that the order in which the images are detected is important. This could be the case if one uses specific lens models (*e.g.* the first one to be detected would always be magnified and the other could be magnified or demagnified). The reversing is also used for testing and validation in order to take into account in the performance the possibly different predictions of the network for a pair and its reversed version.

The strategy for training is to use data with various SNRs in the training set. The SNR of each pair was drawn uniformly in $[5, 20]$ to have low wave amplitudes, but also larger ones that may help the network recognise what it has to learn.

4.5.4 Results

The training was repeated independently $n = 19$ times to take into account the random initialisation of the network. The mean values and error bars that will be shown in the following refer to this difference in training. They are computed as the errors (at 1σ) and averages over the results of the n networks. Note that the error bars always refer to uncertainty with respect to the y -axis quantity.




Figure 4.11: ROC curves for the classifier with different SNR ranges.

The performance of the network will be analysed through different quantities. First, we evaluate the performance using the Receiver Operating Characteristics (ROC) curves. These show the sensitivity, *i.e.* the number of detected lensed events over the total number of lensed events (hence the fraction of lensed events that are detected), as a function of the FPP. In other words, they indicate which proportion of lensed events we can hope to identify

for a given FPP. The ROC plots are shown in Fig. 4.11, where a curve is shown for different SNR ranges. The blue curve represents the range of SNRs on which the network was trained. This curve is the one we need to pay attention to if we try to evaluate the global average performance of the network. The two other curves represent sub-threshold events for the red one, and super-threshold ones for the green curve. This limit distinguishes events that are considered to be detected with enough confidence from those that are not and is set to an SNR of about 8 [67]. Note however that this threshold changes to a combined SNR of 12 when considering several detectors [67]. For events above the detection threshold, we can see that one can reach an FPP of 10^{-2} while preserving a sensitivity above 0.9. Predictions on sub-threshold events can hardly be trusted, or one could set a low FPP knowing that very few lensed events will be identified. On the other hand, the performance on the super-threshold events is better, which is expected given the larger SNRs.




Figure 4.12: ROC plots for the classifier at different SNRs

It can also be interesting to see how the classifier performs using data with a single SNR. This is shown in Fig. 4.12. One can see that at low SNRs, 5 here, the classifier performs poorly, while it gets better and better as the SNR increases. In other words, one can identify more lensed events (greater sensitivity), while having a relatively low probability of wrongly classifying unlensed waves. The difference with an SNR of 8 is significant. All the higher SNR curves converge to a sensitivity of 1 a bit before an FPP of 0.1. This value is very large and is definitely not satisfactory in our case, since we need to minimise the FPP. We can also see that the network performs almost perfectly for an SNR of 30. For these higher SNRs, the sensitivity remains acceptable for a satisfactory FPP.

As mentioned, it may be interesting to see how the classifier performs on individual SNRs. To this end, we can look at other quantities than the sensitivity, such as the accuracy (proportion of correctly classified pairs). It is shown in Fig. 4.13. We can see that the

accuracy keeps on increasing with the SNR, even on SNRs that were not included in the training (25 here). At an SNR of 1, the noise is too large to hope for a good classifier. We can notice that the accuracy crosses the 90% at an SNR of about 7, and at 10, it already reaches an accuracy of 97%. We can then see that the classifier is able to perform relatively well at relatively low optimal SNRs. We can also see that there is little variation with training, given the relatively small error bars.




Figure 4.13: Accuracy as a function of the optimal SNR.

It was mentioned that some waves with close parameters can look alike and might then be wrongly classified as lensed. In fact, it is predicted that the number of these pairs will dominate the number of lensed pairs for a sufficiently large number of detected events [69]. To test the ability of the network to discriminate such events, the FPP was computed on unlensed waves with close chirp masses. The tests were done on 2×10^4 pairs for each relative difference $\Delta M_c/M_c$, which are taken between 1 and 40%. The events generated with the chirp masses M_c and $M_c \pm \Delta M_c$ were given the same mass ratio. The FPP as a function of this relative difference is shown in Fig. 4.14. The FPP vanishes quickly for higher SNRs, at a difference of 5 to 10%. The blue curve represents the average performance over different SNRs and its lower quality is explained by a much poorer performance for lower SNRs (such as 5). It might thus be pessimistic. If we rather trust the SNRs between 10 to 20, the network seems to be able to discriminate waves with large confidence for a relative difference in chirp mass of about 10 to 20%. Note that for this test, the reversed versions of the pairs were not used, as they are irrelevant to the comparison.

Another point to analyse is the usefulness of using a pair and the same one but reversing the order of the event for training. To see how it changes the performance of the network, the same architecture was trained without using the second pair, *i.e.* always with the magnified wave first and the demagnified one second. To account for the uncertainty on weight initialisation, 10 networks were trained and the error bars still refer to this uncertainty. These networks were tested on a test set for which a pair and its reversed version are used. The results are shown in Fig. 4.15. The quantity that is reported is the number of predictions that are different for a pair and the reversed one over the total number of such ‘pairs of

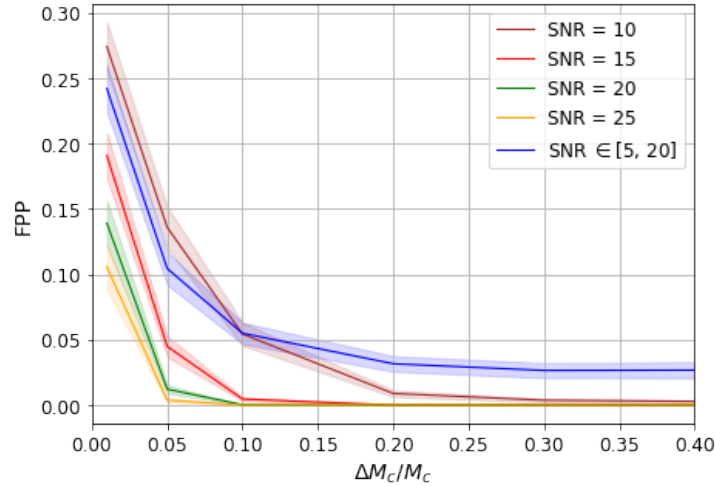


Figure 4.14: FPP for pairs of unlensed events that have a close chirp mass and at different SNRs.

pairs’. We can see that about 50% of the predictions are different if the network is not trained on reversed pairs, while this proportion is much smaller in the other case, at least for reasonable SNRs. At very low SNRs, one still reaches 20%, which is not negligible.

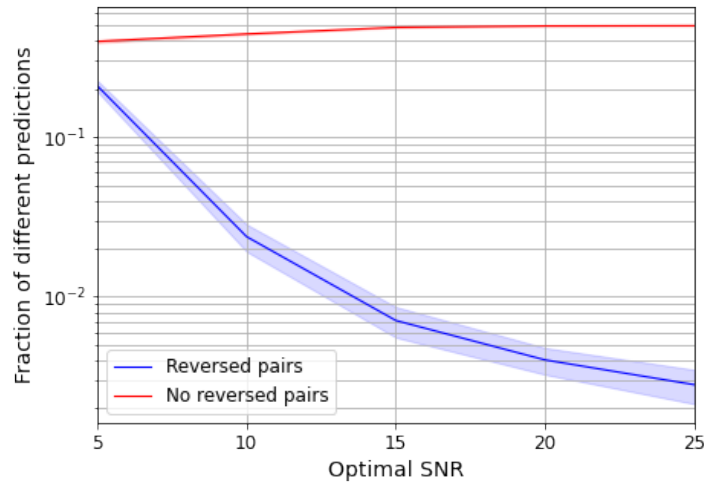


Figure 4.15: Comparison of two training strategies with respect to the order of the events in the pair (magnified or demagnified wave first). For the blue curves, the pair and its reversed version are both used for training. For the red curve, the reversed pair is not used.

Finally, the performance of the classifier on the range of SNRs on which it was trained is shown in Tab. 4.2. The value of the FPP seems reasonable for a first analysis (even within the error bars), and the sensitivity is relatively large as well, which suggests that the classifier is able to correctly recognise lensed waves.

Accuracy	0.963 ± 0.0031
Sensitivity	0.9709 ± 0.0044
FPP	0.045 ± 0.0047

Table 4.2: Performance of the classifier on data taken with a uniform distribution of SNRs in $[5, 20]$.

4.5.5 Discussion

There are several points to discuss in the results. First, regarding the ROC curves, one can choose a threshold to obtain a given FPP. Given Fig. 4.12, we see that we can decrease the value of this parameter with respect to the one shown in Tab. 4.2 at the cost of decreasing the sensitivity. This thus depends on the tolerance we have for the sensitivity. In this case, the curves decrease quite fast with decreasing FPP, so that an FPP of 0.01 leads to a sensitivity a bit larger than 0.8, meaning that about 80% of the lensed events will be detected. If we focus on super-threshold events, we can reach a sensitivity of about 95% for the FPP of 0.01. This seems quite reasonable, in fact, for a first analysis of the data, if it allows one to reduce the number of false positives that would be passed to another algorithm that would confirm the detections. If one rather thinks that one should detect all lensed events at any cost, *i.e.* have a sensitivity of about 1, one can then estimate the corresponding FPP with the ROC curves. From Fig. 4.11, one gets an FPP of about 0.1 for a sensitivity of ~ 1 on average, while for super-threshold events, the FPP is roughly 0.03. In other words, one could divide the number of candidates by a factor of 33 while keeping all the truly lensed events. The conclusion remains the same, the performance also seems encouraging for a first analysis.

It is very interesting to see that the accuracy and sensitivity both increase and that the FPP decreases with the SNR, as seen in the ROC plots and in Fig. 4.13. Given what we see in Fig. 4.10, this seems natural, since the waves are much more visible. But during training, the network was not exposed to data with such significant wave amplitude as obtained with an SNR of 30. This suggests that the network has indeed learned to recognise lensed waves, since it does classify them very well when the waves are more visible.

When comparing the curve with SNRs in $[5, 20]$ in the ROC plots to the ones obtained in [62], the performance seems to be comparable to the model presented in that work. However, I would remain extremely cautious in this comparison, since the data that are used are very different. Moreover, the FPP is less accurate in this work, because in [62], they use a much higher number of unlensed events (with respect to the number of lensed ones), which allows one to better constrain the FPP. They also use 3 detectors, whereas only one is used here. We can nevertheless say that the results are encouraging and that it is worth pursuing the study of the analysis of the time series. If one takes the Morse phase into account, though, it may be that the network does not perform as well, since the waveforms could look different from one another, although being the same but shifted. This phase has however no effect on Q-transforms. One could try to train 3 networks, one on each possible Morse phase.

Another limitation is due to the resemblance of waves that are not lensed, though it

is not clear which proportion of these are found in the data sets. These conditions are however quite realistic, since, as mentioned, many of such pairs are expected to be detected. Therefore, one may be concerned about the limited performance of this network as shown in Fig. 4.14. It is however limited only for SNRs below 15, for which the network correctly identifies most of the unlensed pairs from a relative difference in chirp mass of about 10%, which seems reasonable compared to the accuracy of its measurement. It should be possible to improve the performance by including such pairs in the training data and increasing the sample rate might help as well, since one then increases the number of points on which both waves are compared. In any case, the best performance for this particular task would still be limited by the accuracy of the detectors.

Then, we see in Fig. 4.15 that the chosen training strategy allows one to improve the performance when the order of the events (magnified or demagnified first) does not matter, as expected. On the other hand, there is still a significant difference in predictions at low SNRs. Training without the reversed pairs would have been a limitation to the performance of the network, since it was tested on pairs and reversed pairs. It was however found that testing with the reversed pairs does not change much the reported performance of the network trained with such pairs. Another possible strategy for training is to reverse the order of events in pairs randomly, without using twice the same pair contrarily to what was done here. This could also be applied to the test set.

As there has not been any positive detection so far, one cannot test the performance of the network on real data. It would however be interesting to compare the predictions of this model on the observations to the ones of other techniques.

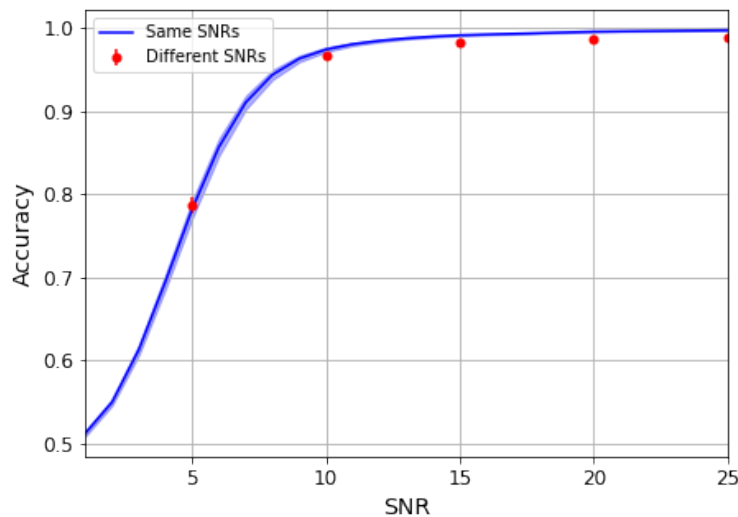


Figure 4.16: Comparison between the accuracy when the amplitudes of the waves of the unlensed pairs are not scaled (blue curve) with the one when they are scaled differently (red dots).

We may wonder which features the CNN recognises in order to classify a pair as lensed or unlensed. Since all lensed events it sees have quite different amplitudes because of the magnification factor, while all unlensed events have a more or less equal amplitude, it might be that the network relies on this difference to classify the pairs. This feature is however

totally irrelevant since, in practice, unlensed waves may have very different amplitudes and SNRs. To see whether this criterion is used by the network, another test was made where the amplitude of each wave of an unlensed pair is multiplied by a different factor. In fact, the multiplicative factors are drawn uniformly from the same ranges as for the magnification factors of the lensed events. The results are shown in Fig. 4.16, where the curve is the same as in Fig. 4.13, while the red dots represents the performance on the test set when the amplitudes of the waves of the unlensed pairs are scaled differently. We can see that the performance of the network is slightly better when the amplitudes of the unlensed events are not changed. If the network relied on the criterion that lensed events display different amplitudes while unlensed ones do not, then the performance would decrease much more than that. One can then safely conclude that this irrelevant criterion is not used by the network, at least not as the main one. This again supports the idea that the network has learned what it must recognise and probably identifies unlensed events through a difference in their waveforms.

Chapter 5

Interests and future prospects

We shall now review what are the interests in detecting lensed gravitational waves, justifying all the work done in the last section. We shall also mention a few points for continuing the work and improving the network that was designed.

5.1 Advantages of the identification of lensed waves

There are several interests in identifying lensed gravitational waves. The first one is that the amplification of the wave allows one to observe more distant sources that would not be visible otherwise. Then, it will allow one to locate the sources better. For example, if the gravitational wave is lensed, its host galaxy is lensed as well in the electromagnetic spectrum. As a result, it may be easier to locate the source of the wave and find the host galaxy [70]. Also, each image can be seen as a new set of detectors for the same source, since the interferometers will not be at the same position and with the same orientation for each image because of the motion of the Earth and the time delay between images. This also allows one to constrain the position of the source better [70]. Contrarily to the multimessenger observation of GW170817 (see Section 4.3.2), one would not observe light from the merger itself, but only from the galaxy. One could still measure the Hubble constant from the redshift of the host galaxy and the distance estimated from the lensed event. It was also proposed that the speed of gravitational waves can be computed from a lensed wave and its electromagnetic counterpart [71] (though already measured from GW170817), and that the Hubble constant can be better constrained with these than with only electromagnetic signals [72]. Also, lensing could be used to test general relativity, but also theories beyond it [73]. Finally, similarly to electromagnetic lensing, it can be used to probe lenses and study their distribution and properties. For example, [74] explores the use of strong lensing to this end. It is also noticed that, contrarily to light, gravitational waves are not submitted to dust extinction. Other features that can be searched for in the lens are small-scale objects, such as dark matter clumps in the halo of a galaxy [75].

It thus appears that lensed waves have a great scientific value. The measurement of the Hubble constant may be particularly interesting, given the current tension in cosmology, with the two values of the Hubble constant that seem incompatible, though relatively close to each other. Also, it is important to test general relativity in different ways to see if it

is still accurate and, if not, find where it might be wrong. Testing other theories may also allow one to discard them.

5.2 Problems with the misidentification of lensed waves

Another important point is that the parameters of the source estimated from the lensed waveform are modified with respect to the true ones because of lensing. In particular, the masses and luminosity distance are affected by this phenomenon [76]. This distance can be estimated from the strain, given that the amplitude decreases linearly with distance [77]. Since the lensing increases or decreases the amplitude, one would then respectively under- and over-estimate the luminosity distance. In the geometrical optics approximation, the luminosity distance is divided by $\sqrt{|\mu|}$, with μ the magnification [12]. There is another quantity that is indirectly influenced by lensing. This is a mass quantity called the *chirp mass*, which is defined as [77]

$$M_c = \frac{(m_1 m_2)^{3/5}}{(m_1 + m_2)^{1/5}}, \quad (5.1)$$

where m_1 and m_2 are the masses of the two compact objects. This parameter appears at the first order of the post-Newtonian expansion of the waveform (corrections of the order v/c to the Newtonian theory). In this context, it can be expressed as [77]

$$M_c \propto \left(f^{-11} \dot{f}^3 \right)^{1/5}, \quad (5.2)$$

where f is the frequency of the gravitational wave and \dot{f} is its time derivative. As a result, the chirp mass that we observe is different from the one at the source, given that the wave is redshifted. At a given redshift, we have [76]

$$M_c(z) = (1 + z) M_{c,0}, \quad (5.3)$$

where $M_{c,0}$ is the chirp mass at the source as presented in Eq. (5.2). Therefore, the mass we observe on Earth is $M_c(z_s)$, with z_s the redshift of the source. In fact, when the wave is lensed without us knowing, we may find a chirp mass $\tilde{M}_c = \tilde{M}_{c,0}(1 + \tilde{z})$, while if it is not lensed we find $M_c = M_{c,0}(1 + z)$. Since the frequency of the wave (and its evolution) is not modified by strong lensing, the observed redshifted chirp mass is the same, *i.e.* $\tilde{M}_c = M_c$, so that [78]

$$\tilde{M}_{c,0}(1 + \tilde{z}) = M_{c,0}(1 + z). \quad (5.4)$$

This is the mass-redshift degeneracy, *i.e.* the redshifted chirp mass can correspond to different values of the chirp mass and the redshift. Therefore, because of lensing, the inferred chirp mass $\tilde{M}_{c,0}$ is different from the true one $M_{c,0}$. The redshift \tilde{z} is the one inferred from the luminosity distance in the case where we do not know the wave is lensed [76]. For the luminosity distance, we have seen that it is divided by $\sqrt{\mu}$, so that in this case we write [76, 78]

$$d_L(\tilde{z}) = \frac{d_L(z)}{\sqrt{|\mu|}}, \quad (5.5)$$

where d_L is the luminosity distance (which depends on the redshift). As a result, the observed distance can be greater or lower (according to μ) than the true distance, and the redshift estimation based on this distance is then also affected by lensing. Lensed events at high redshifts and with a low chirp mass may thus be interpreted as unlensed events at lower redshifts and with a high chirp mass [12, 78].

As a consequence, if some of the events that have been detected so far are lensed without us being able to tell, the population of black holes estimated from all the events may well be biased. What is more, if we cannot tell that two or more images represent the same source in the case of strong lensing, then we consider nonexistent sources in the population. This problem is of great importance, since studying the distribution of different parameters in the population may help consider different formation scenarios of binary black hole systems [76]. For example, several higher-mass black holes were discovered through gravitational waves, heavier than those discovered in the electromagnetic bands [66]. Among the possible explanations, there is the hierarchical merger scenario, *i.e.* successive mergers leading to heavier and heavier black holes [79]. Another possibility is that some events are lensed, leading to larger observed masses than the real ones. If some suggest that a significant proportion of higher redshift lensed events should be detected and that some of the first observing run are lensed [80], others find that the probability of detected events to be strongly lensed, or multiply-imaged, is very small at the detector sensitivity during the first run [81]. Some studies estimated the lensing rate to be of the order of 0.06 lensed event per year or about 5 at LIGO design sensitivity [82], and around 1 per year for aLIGO at design sensitivity for [83]. The latter also predicted that the rate should increase to about 40 to 80 lensed events per year for the Einstein telescope. Given those estimations, the probability that a lensed event has already been detected is very low, though not null. It can be stressed that if a source is multiply-imaged and one of the images is detected, it might be possible that the other image is demagnified and not detectable at the detector sensitivity. Contrarily to longer continuous signals, gravitational waves from mergers are transients, so if an image is not detected, it is lost forever. So, in that case, it would be complicated to identify the detected image as being a lensed wave.

To sum up, lensing affects the observed redshift and chirp mass of the source, which can affect the mass distribution if a lensed wave is not identified as such, and add fictitious sources if there are unidentified multiply-imaged sources. This effect makes it more difficult to identify the origin and nature of the binary system and test formation scenarios. It could explain the large observed masses compared to those of electromagnetic binaries, but it is not the only possible explanation. It also seems unlikely that most of the waves originating from high-mass systems are lensed given the predicted lensing rate and the number of events that have currently been detected.

5.3 Designing a better network

There are several possible improvements that can be made to the model that was proposed. First, one can use all three detectors together instead of projecting the wave on only one of them. This should improve the performance, since the combined SNR would

be greater given the addition of the information on the wave from each detector. Another improvement would be to consider the sky localisation of the events. Two lensed images come from the same position (with the accuracy of the detectors), so that skymaps may help discriminate easily lensed pairs. As already mentioned, there could still be unlensed events that come from very close positions in the sky, so that skymaps alone are not sufficient.

Then, the model can be improved by using more realistic data. A first step could be to vary the sky localisation and the detection time, which takes into account the rotation of the Earth and thus leads to a different projection of the wave. One could also use a given lens model and generate time delays and magnification factors according to this model. The former can be inferred from the time of the detection and then be used as an input to the network. One could also consider a population of lenses and wave sources, although, for the latter at least, I believe we do not have detected enough events yet to have a very accurate representation of the population. Finally, one could use the true noise rather than an idealised one.

Another interesting lead to follow is to leverage the results of deep learning applied to the detection of gravitational waves. One could use two parallel CNNs, each being one designed and pretrained for detection, and pass the outputs of these networks through fully-connected layers. It is however not guaranteed that this will lead to better performance, as the network might not need to have a very accurate representation of the waveform to find differences or similarities between the two images. Indeed, there might be patterns resulting from the combination of the images that are different from those obtained by analysing the waves separately.

Also, it could be interesting to try different learning methods, such as curriculum learning, *i.e.* progressively increase the amplitude of the noise during training. Given the importance of noise in this problem, it might lead to better results. One could again use the results of gravitational-wave detection, for which several training strategies have been studied [84]. There is no guarantee that the best training strategies are the same in our case, but it could be a good starting point given the similarities between the two problems.

Finally, the last improvement could be to change the architecture of the network. Deep learning is a fast-evolving field and new architectures emerge and outperform or, at least, approach the performance of previously state-of-the-art models. For example, transformers [85] were initially designed for translation and were shown to perform very well on image recognition [86]. Therefore, I believe that different models may perform better.

Conclusion

In this work, we have thoroughly explored the phenomenon of gravitational lensing in both the electromagnetic and gravitational wave cases. In particular, we have seen that only the amplitude of the gravitational wave is changed in the geometrical optics approximation and that a phase change is induced as well. This approximation is valid when the wavelength of the wave is much smaller than the Schwarzschild radius of the lens, *i.e.* typically when the lens is a galaxy or a cluster of galaxies. We have also investigated the wave effects that are negligible in the previous approximation. This was done by solving the wave equation in the metric that describes a Minkowski spacetime perturbed by a point-mass lens.

Though these effects of microlensing can be probed using deep learning, we have rather investigated the case of strong macrolensing. The main contribution of this work is the proposition of a neural network architecture to identify pairs of images in the coalescence events that are detected with the interferometers. A different approach than existing works was used, consisting in using the time series representation rather than the spectrogram of the waves. The goal of the model was to show that such a representation allows one to get good results in the identification of lensed waves in a simple case, or at least allows one to discard unlensed events. It is considered that the model achieves performance good enough to consider this objective as fulfilled. In particular, the network shows promising performance on average, in particular at larger SNRs, and seems to understand what it must identify. Several limitations were nevertheless pointed out. The network is not very good at distinguishing unlensed waves that have close parameters. The performance at low SNRs is not good either, but this should not be too much of a concern, for waves are considered to be detected if they have a given minimum SNR.

Although there is no guarantee that this model will outperform existing techniques in more general conditions, we can conclude that the use of deep learning on the time series is an interesting lead to follow. It can be further studied to confirm these results and improve them. Among other things, it can be generalised to multiple detectors and more realistic data. It can also be adapted to the cases where one looks for triplets or quadruplets.

Appendix A

Derivation of relations in electromagnetic lensing

In this section, we shall derive the expression of several quantities presented in Chapter 2. Just as that chapter, it is largely based on [22], [23] and [24].

A.1 Time delay and lensing potential

To set the context again, we consider the fact that the ray takes more time to reach the observer because of its deflection and different metric with respect to an unlensed case. The geometry of the system is represented in Fig. A.1




Figure A.1: Sketch of the geometry of the problem. Dimensions are not to scale, angles are exaggerated compared to distances.

To find the delay that is introduced, we start from the metric (2.11) that reads

$$ds^2 = -(1 + 2U) dt^2 + (1 - 2U) (d\vec{r})^2. \quad (\text{A.1})$$

Since $ds^2 = 0$ for a light ray, we get

$$dt = \sqrt{\frac{1-2U}{1+2U}} |d\vec{r}| = \sqrt{\frac{1-2U}{1+2U}} dl, \quad (\text{A.2})$$

where dl is the arc length in Euclidean space. We can then approximate the square root by $1 - 2U + O(U^2)$ and assume that light is emitted at $t = 0$, so that we get, by integration,

$$t \simeq \int (1 - 2U) dl = l - 2 \int U dl, \quad (\text{A.3})$$

which represents the time required to reach the observer. We can see two contributions to this time: a geometrical one and a gravitational one, each representing the geometrical and gravitational time delays, respectively.

For the geometrical time delay, we can easily compute the length of the path using Fig. A.1, assuming first that the distances are Euclidean ones. To distinguish them from angular diameter distances, I shall refer to them with d instead of D . We then use twice Pythagoras's theorem to express the length of the path from S to the lens plane, l_{SL} , and from the lens plane to O , l_{LO} , we obtain

$$l = l_{SL} + l_{LO} = \sqrt{(\vec{\xi} - \vec{\eta})^2 + d_{LS}^2} + \sqrt{(\vec{\xi})^2 + d_L^2} \quad (\text{A.4})$$

$$= d_{LS} \sqrt{1 + \frac{(\vec{\xi} - \vec{\eta})^2}{d_{LS}^2}} + d_L \sqrt{1 + \frac{(\vec{\xi})^2}{d_L^2}}. \quad (\text{A.5})$$

The small-angle approximation and large distances allow us to consider that $(\vec{\xi} - \vec{\eta})^2 \ll d_{LS}^2$ and $(\vec{\xi})^2 \ll d_L^2$. We can therefore approximate the length by

$$l \simeq d_{LS} + \frac{1}{2d_{LS}} (\vec{\xi} - \vec{\eta})^2 + d_L + \frac{1}{2d_L} (\vec{\xi})^2 \quad (\text{A.6})$$

The time delay Δt is defined as the difference in arrival time between the lensed ray and the unlensed one, the latter being defined purely geometrically (as there is no lens), *i.e.*

$$t_u = l_u = \sqrt{d_S^2 + (\vec{\eta})^2} \simeq d_S + \frac{1}{2d_S} (\vec{\eta})^2 \quad (\text{A.7})$$

In our approximation of Euclidean distances, we have $d_S = d_L + d_{LS}$. As a consequence, we can write the time delay as

$$\Delta t = t - t_u = \frac{1}{2d_{LS}} (\vec{\xi} - \vec{\eta})^2 + \frac{1}{2d_L} (\vec{\xi})^2 - \frac{1}{2d_S} (\vec{\eta})^2 - 2 \int U dl. \quad (\text{A.8})$$

We can further develop this expression into

$$\begin{aligned}
\Delta t &= \frac{1}{2} (\vec{\xi})^2 \left(\frac{1}{d_{LS}} + \frac{1}{d_L} \right) + \frac{1}{2} (\vec{\eta})^2 \left(\frac{1}{d_{LS}} - \frac{1}{d_S} \right) - \frac{1}{d_{LS}} \vec{\xi} \cdot \vec{\eta} - 2 \int U dl \\
&= \frac{1}{2} \frac{d_L + d_{LS}}{d_{LS} d_L} (\vec{\xi})^2 + \frac{1}{2} \frac{d_S - d_{LS}}{d_{LS} d_S} (\vec{\eta})^2 - \frac{1}{d_{LS}} \vec{\xi} \cdot \vec{\eta} - 2 \int U dl \\
&= \frac{1}{2} \frac{d_S}{d_{LS} d_L} (\vec{\xi})^2 + \frac{1}{2} \frac{d_L}{d_{LS} d_S} (\vec{\eta})^2 - \frac{1}{d_{LS}} \vec{\xi} \cdot \vec{\eta} - 2 \int U dl \\
&= \frac{1}{2} \frac{d_S}{d_{LS} d_L} \left[(\vec{\xi})^2 + \frac{d_L^2}{d_S^2} (\vec{\eta})^2 - 2 \frac{d_L}{d_S} \vec{\xi} \cdot \vec{\eta} \right] - 2 \int U dl \\
&= \frac{d_S}{d_L d_{LS}} \left[\frac{1}{2} \left(\vec{\xi} - \frac{d_L}{d_S} \vec{\eta} \right)^2 - 2 \frac{d_{LS} d_L}{d_S} \int U dl \right] \tag{A.9}
\end{aligned}$$

If we inject the dimensionless variables as introduced in Eq. (2.6), (2.7) and (2.8), we get:

$$\begin{aligned}
\Delta t &= \frac{d_S}{d_L d_{LS}} \left[\frac{1}{2} \left(\xi_0 \vec{x} - \frac{d_L}{d_S} \xi_0 \frac{d_S}{d_L} \vec{y} \right)^2 - 2 \frac{d_{LS} d_L}{d_S} \int U dl \right] \\
&= \frac{d_S \xi_0^2}{d_L d_{LS}} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - 2 \frac{d_{LS} d_L}{d_S \xi_0^2} \int U dl \right] \\
\Rightarrow \quad \Delta t(\vec{x}) &= \frac{d_S \xi_0^2}{d_L d_{LS}} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}) \right], \tag{A.10}
\end{aligned}$$

where we have defined the dimensionless lensing potential

$$\psi(\vec{x}) = 2 \frac{d_{LS} d_L}{d_S \xi_0^2} \int U(\vec{x}, l) dl. \tag{A.11}$$

As the deflection angle is small, we can approximate the integral over the true path by the integral over the path of an unlensed ray, with the third dimension Z taken along this ray. We can thus write

$$\psi(\vec{x}) = 2 \frac{d_{LS} d_L}{d_S \xi_0^2} \int U(\vec{x}, Z) dZ. \tag{A.12}$$

If we now introduce cosmological considerations, we need to use the FLRW metric perturbed by a lens. It reads [23, 24]

$$ds^2 = -(1 + 2U) dt^2 + a^2(t) (1 - 2U) (d\vec{r})^2. \tag{A.13}$$

If we perform the same manipulations as in the previous case, we get

$$dt = \sqrt{\frac{1 - 2U}{1 + 2U}} a(t) |d\vec{r}|. \tag{A.14}$$

We will look at the time delay at the observer, so that we consider $a(t) \simeq a_0$. Indeed, the typical time for a light ray to reach us is much shorter than the relative variation of $a(t)$ nowadays (the Hubble time H_0^{-1} , of the order of the age of the Universe), so we can make the approximation that $a(t)$ is constant in that time interval. Recalling that $a_0 = 1$, we retrieve the exact same relation as before. There is nevertheless a difference hiding in the distances that are used. Since at the beginning we used the angular diameter distance, we can use the relation (1.39) to express the comoving distance in terms of the angular diameter distance. We will therefore make the following substitutions:

$$d_L \rightarrow D_L (1 + z_L) \quad d_S \rightarrow D_S (1 + z_S) \quad d_{LS} \rightarrow D_{LS} (1 + z_S) \quad (\text{A.15})$$

Note that it is not as straightforward as this, but it is equivalent, in the end, to what is obtained in flat space in a more rigorous development (for such developments, see [23] and [24]). Note also that one needs to consider the redshift of the object that is observed, irrespective of where the distance is measured (*e.g.* from the lens or from the observer). For d_S and d_{LS} , one measures distances to the source, so one needs to consider a at the time the photon leaves the source, *i.e.* a_S . For d_L one measures the distance to the lens, so one uses a_L .

The case of the integral in $\psi(\vec{x})$ might be a bit more tricky to transform. One argument (in [24]) is to say that the delay due to the potential is localised around the lens, as it is fairly negligible at greater distances. This means that the whole integral is significant only near the lens, so that the time delay due to the potential can be roughly estimated at the lens. As a consequence, one can consider the time delay at the lens and ‘propagate’ it in an appropriate manner to the observer. The way to do this was introduced through the cosmological time dilation and its expression is given in Eq. (1.26). In this case, writing z_L the redshift of the lens, Δt_L the time delay at the lens, and with $a_0 = 1$, we have

$$\Delta t = \Delta t_L (1 + z_L), \quad (\text{A.16})$$

since we can consider that the photons come from the lens and are emitted with a time difference Δt_L ; we treat it as if it were a source. As a consequence and given that the potential time delay is caused by $\psi(\vec{x})$, we write the transformation

$$\psi(\vec{x}) \rightarrow (1 + z_L) \psi(\vec{x}). \quad (\text{A.17})$$

We can also note that ξ_0 then transforms as $\xi_0 (1 + z_L)$, since $\theta_0 = \xi_0 d_L \rightarrow \xi_0 (1 + z_L) D_L$. We then apply the transformations (A.15) and (A.17) on the equations (A.10) and (A.12), and we get

$$\begin{aligned} \psi(\vec{x}) &= 2(1 + z_L) \frac{D_{LS} (1 + z_S) D_L (1 + z_L)}{D_S (1 + z_S) \xi_0^2 (1 + z_L)^2} \int U(\vec{x}, Z) dZ, \\ \Delta t(\vec{x}) &= \frac{D_S (1 + z_S) \xi_0^2 (1 + z_L)^2}{D_{LS} (1 + z_S) D_L (1 + z_L)} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}) \right]. \end{aligned}$$

After simplification, we get

$$\psi(\vec{x}) = 2 \frac{D_{LS} D_L}{D_S \xi_0^2} \int U(\vec{x}, Z) dZ, \quad (\text{A.18})$$

$$\Delta t(\vec{x}) = (1 + z_L) \frac{D_S \xi_0^2}{D_{LS} D_L} \left[\frac{1}{2} (\vec{x} - \vec{y})^2 - \psi(\vec{x}) \right]. \quad (\text{A.19})$$

We can retrieve the relations of Section 2.3 if we choose $\xi_0 = D_L$.

A.2 Deflection angle for a point-mass lens

After a few developments, we found the expression of the deflection angle, which is

$$\vec{\alpha}(\vec{x}) = 2 \frac{D_{LS}}{D_S D_L} \int \vec{\nabla}_{\vec{x}} U(\vec{x}, Z) dZ = 2 \frac{D_{LS}}{D_S} \int \vec{\nabla}_{\perp} U(\vec{x}, Z) dZ, \quad (\text{A.20})$$

Now that we have an expression for the deflection angle, we can compute it for a point-mass lens. It is more convenient to work with $\hat{\vec{\alpha}} = \vec{\alpha} D_S / D_{LS}$ and $\vec{\xi}$ rather than \vec{x} . The relation (A.20) can then be written as

$$\hat{\vec{\alpha}}(\vec{\xi}) = 2 \int \vec{\nabla}_{\perp} U(\vec{\xi}, Z) dZ, \quad (\text{A.21})$$

with the approximation $\vec{\nabla}_{\perp} = \vec{\nabla}_{\vec{\xi}}$. We consider a point-mass lens characterised by the potential $U(\vec{\xi}, Z) = -\frac{GM}{|\vec{r}|}$, with $\vec{r} = (\xi_1, \xi_2, Z) = \vec{\xi} + Z \vec{e}_Z$ and $\vec{\xi} = (\xi_1, \xi_2)$. We then have

$$\vec{\nabla}_{\vec{\xi}} U(\vec{\xi}, Z) = \frac{GM}{2} \frac{\vec{\nabla}_{\vec{\xi}} \left((\vec{\xi})^2 + Z^2 \right)}{|\vec{r}|^3} = GM \frac{\vec{\xi}}{|\vec{r}|^3}. \quad (\text{A.22})$$

Thus, we need to compute

$$\hat{\vec{\alpha}}(\vec{\xi}) = 2 \int GM \frac{\vec{\xi}}{|\vec{r}|^3} dZ = 2GM \vec{\xi} \int \frac{1}{(\xi_1^2 + \xi_2^2 + Z^2)^{3/2}} dZ. \quad (\text{A.23})$$

We can use a change of variables $Z / \sqrt{\xi_1^2 + \xi_2^2} = \tan \alpha$, which leads to

$$\int \frac{1}{(\xi_1^2 + \xi_2^2 + Z^2)^{3/2}} dZ = \frac{2}{\xi_1^2 + \xi_2^2} = \frac{2}{|\vec{\xi}|^2}. \quad (\text{A.24})$$

Injecting this result in the definition of the deflection angle, we obtain

$$\hat{\vec{\alpha}}(\vec{\xi}) = 4GM \frac{\vec{\xi}}{|\vec{\xi}|^2}, \quad (\text{A.25})$$

$$|\hat{\vec{\alpha}}(\vec{\xi})| = \frac{4GM}{\xi}, \quad (\text{A.26})$$

with $\xi = |\vec{\xi}|$ the impact parameter. This result can also be obtained from general relativity in the Schwarzschild metric.

A.3 Magnification for a point-mass lens

In this section, we shall briefly derive the expression for the magnification resulting from lensing by a point-mass lens. This is based on [22].

First, let us recall that the solutions of the lens equation $y = x - 1/x$ are

$$x_{\pm} = \frac{1}{2} \left(y \pm \sqrt{y^2 + 4} \right). \quad (\text{A.27})$$

In the case of an axially-symmetric lens, so also for a point-mass lens, the determinant of the Jacobian is expressed as

$$\det A = \frac{y}{x} \frac{\partial y}{\partial x}. \quad (\text{A.28})$$

We thus have, using the lens equation $y = x - 1/x$,

$$\det A = \left(1 - \frac{1}{x^2} \right) \left(1 + \frac{1}{x^2} \right) = 1 - \frac{1}{x^4}, \quad (\text{A.29})$$

and the magnification, which is the inverse of this determinant, reads

$$\mu = \frac{1}{1 - \frac{1}{x^4}} = \frac{x^4}{x^4 - 1}. \quad (\text{A.30})$$

In particular, we can compute the magnification for the images that are solution to the lens equation. To that end, we need first to notice that, after a bit of algebra, we have the following relations:

$$x_{\pm} + \frac{1}{x_{\pm}} = \pm \sqrt{y^2 + 4} \quad (\text{A.31})$$

$$x_{\pm}^2 = \frac{1}{2} \left(y^2 + 2 \pm y \sqrt{y^2 + 4} \right) \quad (\text{A.32})$$

So, the magnifications for the observed images are

$$\begin{aligned} \mu_{\pm} &= \frac{1}{1 - \frac{1}{x_{\pm}^4}} = \frac{1}{\left(1 - \frac{1}{x_{\pm}^2} \right) \left(1 + \frac{1}{x_{\pm}^2} \right)} \\ &= \frac{x_{\pm}^2}{\left(x_{\pm} - \frac{1}{x_{\pm}} \right) \left(x_{\pm} + \frac{1}{x_{\pm}} \right)} \\ &= \frac{1}{2} \frac{y^2 + 2 \pm y \sqrt{y^2 + 4}}{\pm y \sqrt{y^2 + 4}}, \end{aligned} \quad (\text{A.33})$$

where the two relations given above are used in the last step, along with the lens equation $y = x - 1/x$. We can still simplify a bit the expression and we obtain

$$\mu_{\pm} = \frac{1}{2} \pm \frac{y^2 + 2}{2y \sqrt{y^2 + 4}}. \quad (\text{A.34})$$

Appendix B

Derivation of the wave equation

In this section, we shall derive the wave equation with detailed calculations. The developments are based on [26]. As stated in Section 3.1.2, we start from Einstein's equations (1.16). To express these equations, we consider a background metric $\tilde{g}_{\mu\nu}$ with a perturbation $h_{\mu\nu}$, such that

$$g_{\mu\nu} = \tilde{g}_{\mu\nu} + h_{\mu\nu} , \quad (\text{B.1})$$

and

$$g^{\mu\nu} \simeq \tilde{g}^{\mu\nu} - h^{\mu\nu} . \quad (\text{B.2})$$

To first order, $\tilde{g}^{\mu\nu}$ raises the indices of $h_{\mu\nu}$. The background metric is such that it varies on much larger scales than the wavelength of the perturbation $h_{\mu\nu}$ and it solves Einstein's equation in vacuum (see Section 3.1.2).

The next step is to derive the Ricci tensor and scalar to express Einstein's equations. To this end, we must first express the Levi-Civita connection, with its expression shown in equation (1.12). This expression can be linearised, so that we can express the connection in our metric (B.1) as

$$\Gamma^\alpha{}_{\mu\nu} \simeq \tilde{\Gamma}^\alpha{}_{\mu\nu} + \delta\Gamma^\alpha{}_{\mu\nu} , \quad (\text{B.3})$$

where the first term on the right-hand side is the connection defined by the background metric and the second one is due to the presence of the perturbation $h_{\mu\nu}$. Indeed, we get, from (1.12),

$$\begin{aligned} \Gamma^\alpha{}_{\mu\nu} &\simeq \frac{1}{2} (\tilde{g}^{\alpha\beta} - h^{\alpha\beta}) (\partial_\mu \tilde{g}_{\nu\beta} + \partial_\nu \tilde{g}_{\mu\beta} - \partial_\beta \tilde{g}_{\mu\nu} + \partial_\mu h_{\nu\beta} + \partial_\nu h_{\mu\beta} - \partial_\beta h_{\mu\nu}) \\ &= \tilde{\Gamma}^\alpha{}_{\mu\nu} + \frac{1}{2} \tilde{g}^{\alpha\beta} (\partial_\mu h_{\nu\beta} + \partial_\nu h_{\mu\beta} - \partial_\beta h_{\mu\nu}) - \frac{1}{2} h^{\alpha\beta} (\partial_\mu \tilde{g}_{\nu\beta} + \partial_\nu \tilde{g}_{\mu\beta} - \partial_\beta \tilde{g}_{\mu\nu}) + O(h^2) , \\ &= \tilde{\Gamma}^\alpha{}_{\mu\nu} + \delta\Gamma^\alpha{}_{\mu\nu} + O(h^2) , \end{aligned} \quad (\text{B.4})$$

since the terms $h^{\alpha\beta}(\partial_\mu h_{\nu\beta} + \dots)$ are second order terms. We have thus identified the term " $\delta\Gamma$ ":

$$\delta\Gamma^\alpha{}_{\mu\nu} = \frac{1}{2} \tilde{g}^{\alpha\beta} (\partial_\mu h_{\nu\beta} + \partial_\nu h_{\mu\beta} - \partial_\beta h_{\mu\nu}) - \frac{1}{2} h^{\alpha\beta} (\partial_\mu \tilde{g}_{\nu\beta} + \partial_\nu \tilde{g}_{\mu\beta} - \partial_\beta \tilde{g}_{\mu\nu}) . \quad (\text{B.5})$$

To develop this expression further, we need to recall the property (1.11) of the Levi-Civita connection, which implies that (using equation (1.8))

$$\begin{aligned}\nabla_\alpha \tilde{g}_{\mu\nu} &= \partial_\alpha \tilde{g}_{\mu\nu} - \tilde{\Gamma}^\lambda_{\alpha\mu} \tilde{g}_{\nu\lambda} - \tilde{\Gamma}^\lambda_{\alpha\nu} \tilde{g}_{\mu\lambda} = 0 \\ \implies \partial_\alpha \tilde{g}_{\mu\nu} &= \tilde{\Gamma}^\lambda_{\alpha\mu} \tilde{g}_{\nu\lambda} + \tilde{\Gamma}^\lambda_{\alpha\nu} \tilde{g}_{\mu\lambda},\end{aligned}\tag{B.6}$$

$$\implies \partial_\alpha \tilde{g}^{\mu\nu} = -\tilde{\Gamma}^\mu_{\alpha\lambda} \tilde{g}^{\nu\lambda} - \tilde{\Gamma}^\nu_{\alpha\lambda} \tilde{g}^{\mu\lambda}.\tag{B.7}$$

We can use this to simplify the second term on the right-hand side of Eq. (B.5). Taking into account the symmetry property (1.10) of the Levi-Civita connection (and the symmetry of the metric), we obtain

$$\begin{aligned}\partial_\mu \tilde{g}_{\nu\beta} + \partial_\nu \tilde{g}_{\mu\beta} - \partial_\beta \tilde{g}_{\mu\nu} &= \tilde{\Gamma}^\lambda_{\mu\nu} \tilde{g}_{\beta\lambda} + \tilde{\Gamma}^\lambda_{\mu\beta} \tilde{g}_{\nu\lambda} + \tilde{\Gamma}^\lambda_{\nu\mu} \tilde{g}_{\beta\lambda} + \tilde{\Gamma}^\lambda_{\nu\beta} \tilde{g}_{\mu\lambda} - \tilde{\Gamma}^\lambda_{\beta\mu} \tilde{g}_{\nu\lambda} - \tilde{\Gamma}^\lambda_{\beta\nu} \tilde{g}_{\mu\lambda} \\ &= \tilde{g}_{\beta\lambda} \left(\tilde{\Gamma}^\lambda_{\mu\nu} + \tilde{\Gamma}^\lambda_{\nu\mu} \right) + \tilde{g}_{\nu\lambda} \left(\tilde{\Gamma}^\lambda_{\mu\beta} - \tilde{\Gamma}^\lambda_{\beta\mu} \right) + \tilde{g}_{\mu\lambda} \left(\tilde{\Gamma}^\lambda_{\nu\beta} - \tilde{\Gamma}^\lambda_{\beta\nu} \right) \\ &= 2 \tilde{g}_{\beta\lambda} \tilde{\Gamma}^\lambda_{\mu\nu}.\end{aligned}\tag{B.8}$$

The full term thus reads

$$\frac{1}{2} h^{\alpha\beta} (\partial_\mu \tilde{g}_{\nu\beta} + \partial_\nu \tilde{g}_{\mu\beta} - \partial_\beta \tilde{g}_{\mu\nu}) = h^{\alpha\beta} \tilde{g}_{\beta\lambda} \tilde{\Gamma}^\lambda_{\mu\nu} \simeq h_\lambda^\alpha \tilde{\Gamma}^\lambda_{\mu\nu}.\tag{B.9}$$

We can then use the fact that $\tilde{g}^{\mu\nu}$ raises indices to first order and the relation (B.7) to get

$$\begin{aligned}\tilde{g}^{\alpha\beta} (\partial_\mu h_{\nu\beta} + \partial_\nu h_{\mu\beta} - \partial_\beta h_{\mu\nu}) &\simeq (\partial_\mu h_\nu^\alpha - h_{\nu\beta} \partial_\mu \tilde{g}^{\beta\lambda} + \partial_\nu h_\mu^\alpha - h_{\mu\beta} \partial_\nu \tilde{g}^{\beta\lambda} - \tilde{g}^{\alpha\beta} \partial_\beta h_{\mu\nu}) \\ &= \partial_\mu h_\nu^\alpha + h_{\nu\beta} \tilde{\Gamma}^\alpha_{\mu\lambda} \tilde{g}^{\beta\lambda} + h_{\nu\beta} \tilde{\Gamma}^\beta_{\mu\lambda} \tilde{g}^{\alpha\lambda} + \partial_\nu h_\mu^\alpha + h_{\mu\beta} \tilde{\Gamma}^\alpha_{\nu\lambda} \tilde{g}^{\beta\lambda} + h_{\mu\beta} \tilde{\Gamma}^\beta_{\nu\lambda} \tilde{g}^{\alpha\lambda} - \tilde{g}^{\alpha\beta} \partial_\beta h_{\mu\nu} \\ &= \partial_\mu h_\nu^\alpha + h_\nu^\lambda \tilde{\Gamma}^\alpha_{\mu\lambda} + h_{\nu\beta} \tilde{\Gamma}^\beta_{\mu\lambda} \tilde{g}^{\alpha\lambda} + \partial_\nu h_\mu^\alpha + h_\mu^\lambda \tilde{\Gamma}^\alpha_{\nu\lambda} + h_{\mu\beta} \tilde{\Gamma}^\beta_{\nu\lambda} \tilde{g}^{\alpha\lambda} - \tilde{g}^{\alpha\beta} \partial_\beta h_{\mu\nu}\end{aligned}\tag{B.10}$$

We then put everything together in equation (B.5) and arrange a bit the terms to get

$$\begin{aligned}\delta \Gamma^\alpha_{\mu\nu} &= \frac{1}{2} \left[\left(\partial_\mu h_\nu^\alpha - h_\lambda^\alpha \tilde{\Gamma}^\lambda_{\mu\nu} + h_\nu^\lambda \tilde{\Gamma}^\alpha_{\mu\lambda} \right) + \left(\partial_\nu h_\mu^\alpha - h_\lambda^\alpha \tilde{\Gamma}^\lambda_{\nu\mu} + h_\mu^\lambda \tilde{\Gamma}^\alpha_{\nu\lambda} \right) \right. \\ &\quad \left. - g^{\alpha\lambda} \left(\partial_\lambda h_{\mu\nu} - h_{\mu\beta} \tilde{\Gamma}^\beta_{\nu\lambda} - h_{\nu\beta} \tilde{\Gamma}^\beta_{\mu\lambda} \right) \right] \\ &= \frac{1}{2} \left(\tilde{\nabla}_\mu h_\nu^\alpha + \tilde{\nabla}_\nu h_\mu^\alpha - \tilde{\nabla}^\alpha h_{\mu\nu} \right),\end{aligned}\tag{B.11}$$

where we used the definition of the covariant derivative (Eq. (1.7) and (1.8)) and $\tilde{\nabla}^\alpha = \tilde{g}^{\alpha\lambda} \tilde{\nabla}_\lambda$.

Similarly to the connection, we can linearise the expression of the Riemann tensor to write

$$R^\alpha_{\beta\mu\nu} \simeq \tilde{R}^\alpha_{\beta\mu\nu} + \delta R^\alpha_{\beta\mu\nu}\tag{B.12}$$

If we use the expression (1.13) of the Riemann tensor, noting that $\delta \Gamma^\alpha_{\mu\nu}$ is $O(h)$, and develop $\Gamma^\alpha_{\mu\nu}$ as in Eq. (B.3), we have:

$$\begin{aligned}
R^\alpha_{\beta\mu\nu} &= \partial_\mu \Gamma^\alpha_{\nu\beta} - \partial_\nu \Gamma^\alpha_{\mu\beta} + \Gamma^\alpha_{\mu\lambda} \Gamma^\lambda_{\nu\beta} - \Gamma^\alpha_{\nu\lambda} \Gamma^\lambda_{\mu\beta} \\
&= \tilde{R}^\alpha_{\beta\mu\nu} + \partial_\mu \delta \Gamma^\alpha_{\nu\beta} - \partial_\nu \delta \Gamma^\alpha_{\mu\beta} + \tilde{\Gamma}^\alpha_{\mu\lambda} \delta \Gamma^\lambda_{\nu\beta} + \tilde{\Gamma}^\lambda_{\nu\beta} \delta \Gamma^\alpha_{\mu\lambda} \\
&\quad - \tilde{\Gamma}^\alpha_{\nu\lambda} \delta \Gamma^\lambda_{\mu\beta} - \tilde{\Gamma}^\lambda_{\mu\beta} \delta \Gamma^\alpha_{\nu\lambda} + O(h^2), \tag{B.13}
\end{aligned}$$

since again the terms " $\delta\Gamma \delta\Gamma$ " are second order terms. We can then identify the $\delta R^\alpha_{\beta\mu\nu}$ part. The presence of a derivative of $\delta \Gamma^\alpha_{\mu\nu}$ with its multiplication with the background connection allows one to introduce the covariant derivative in the expression. A term is however missing to get the full covariant derivative so that we get

$$\begin{aligned}
\delta R^\alpha_{\beta\mu\nu} &= \partial_\mu \delta \Gamma^\alpha_{\nu\beta} - \partial_\nu \delta \Gamma^\alpha_{\mu\beta} + \tilde{\Gamma}^\alpha_{\mu\lambda} \delta \Gamma^\lambda_{\nu\beta} + \tilde{\Gamma}^\lambda_{\nu\beta} \delta \Gamma^\alpha_{\mu\lambda} - \tilde{\Gamma}^\alpha_{\nu\lambda} \delta \Gamma^\lambda_{\mu\beta} - \tilde{\Gamma}^\lambda_{\mu\beta} \delta \Gamma^\alpha_{\nu\lambda} \\
&= \left(\partial_\mu \delta \Gamma^\alpha_{\nu\beta} + \tilde{\Gamma}^\alpha_{\mu\lambda} \delta \Gamma^\lambda_{\nu\beta} - \tilde{\Gamma}^\lambda_{\mu\beta} \delta \Gamma^\alpha_{\nu\lambda} \right) - \left(\partial_\nu \delta \Gamma^\alpha_{\mu\beta} + \tilde{\Gamma}^\alpha_{\nu\lambda} \delta \Gamma^\lambda_{\mu\beta} - \tilde{\Gamma}^\lambda_{\nu\beta} \delta \Gamma^\alpha_{\mu\lambda} \right) \\
&= \left(\tilde{\nabla}_\mu (\delta \Gamma^\alpha_{\nu\beta}) - \tilde{\Gamma}^\lambda_{\mu\nu} \delta \Gamma^\alpha_{\lambda\beta} \right) - \left(\tilde{\nabla}_\nu (\delta \Gamma^\alpha_{\mu\beta}) - \tilde{\Gamma}^\lambda_{\nu\mu} \delta \Gamma^\alpha_{\lambda\beta} \right) \\
&= \tilde{\nabla}_\mu (\delta \Gamma^\alpha_{\nu\beta}) - \tilde{\nabla}_\nu (\delta \Gamma^\alpha_{\mu\beta}), \tag{B.14}
\end{aligned}$$

where we used the symmetry property of the connection in the last step. We can now inject the expression of the $\delta\Gamma$ terms found in Eq. (B.11) and we finally end up with

$$\delta R^\alpha_{\beta\mu\nu} = \frac{1}{2} \left(\tilde{\nabla}_\mu \tilde{\nabla}_\nu h_\beta^\alpha - \tilde{\nabla}_\nu \tilde{\nabla}_\mu h_\beta^\alpha + \tilde{\nabla}_\mu \tilde{\nabla}_\beta h_\nu^\alpha - \tilde{\nabla}_\nu \tilde{\nabla}_\beta h_\mu^\alpha - \tilde{\nabla}_\mu \tilde{\nabla}^\alpha h_{\nu\beta} + \tilde{\nabla}_\nu \tilde{\nabla}^\alpha h_{\mu\beta} \right). \tag{B.15}$$

We can then easily compute the Ricci tensor and scalar from Eq. (1.14) and (1.15). For the Ricci tensor, we can write

$$R_{\mu\nu} = R^\alpha_{\mu\alpha\nu} = \tilde{R}_{\mu\nu} + \delta R^\alpha_{\mu\alpha\nu} = \tilde{R}_{\mu\nu} + \delta R_{\mu\nu}. \tag{B.16}$$

We can also notice that $\tilde{\nabla}^\alpha h_{\mu\alpha} = \tilde{\nabla}_\alpha h_\mu^\alpha$, given the properties of the Levi-Civita connection, especially Eq. (1.11). We can then calculate

$$\begin{aligned}
2 \delta R_{\mu\nu} &= \tilde{\nabla}_\alpha \tilde{\nabla}_\nu h_\mu^\alpha - \tilde{\nabla}_\nu \tilde{\nabla}_\alpha h_\mu^\alpha + \tilde{\nabla}_\alpha \tilde{\nabla}_\mu h_\nu^\alpha - \tilde{\nabla}_\nu \tilde{\nabla}_\mu h_\alpha^\alpha - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h_{\nu\mu} + \tilde{\nabla}_\nu \tilde{\nabla}^\alpha h_{\alpha\mu} \\
&= \tilde{\nabla}_\alpha \tilde{\nabla}_\nu h_\mu^\alpha + \tilde{\nabla}_\alpha \tilde{\nabla}_\mu h_\nu^\alpha - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h_{\nu\mu} - \tilde{\nabla}_\nu \tilde{\nabla}_\mu h, \tag{B.17}
\end{aligned}$$

where we define $h_\alpha^\alpha = h = \text{Tr}[h_{\alpha\beta}]$.

Finally, the last step before expressing the Einstein tensor is to write the Ricci scalar. We have

$$\delta R = g^{\mu\nu} \delta R_{\mu\nu} = \tilde{g}^{\mu\nu} \delta R_{\mu\nu} + O(h^2), \tag{B.18}$$

since $\delta R_{\mu\nu}$ is $O(h)$. The $\tilde{g}^{\mu\nu}$ can enter inside the covariant derivatives (given the property (1.11)) and can raise the indices of h . So, developing the δR term yields

$$\begin{aligned}\delta R &= \tilde{g}^{\mu\nu} \delta R_{\mu\nu} \\ &= \frac{1}{2} \left(\tilde{\nabla}_\alpha \tilde{\nabla}_\nu h^{\nu\alpha} + \tilde{\nabla}_\alpha \tilde{\nabla}_\mu h^{\mu\alpha} - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h_\nu{}^\nu - \tilde{\nabla}_\nu \tilde{\nabla}^\nu h \right) \\ &= \tilde{\nabla}_\alpha \tilde{\nabla}_\beta h^{\beta\alpha} - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h,\end{aligned}\tag{B.19}$$

changing the names of the dummy indices μ and ν .

Now that we have computed all the different components of Einstein's equations, we can express the Einstein tensor as defined in Eq. (1.16). As δR is $O(h)$, we have

$$\begin{aligned}G_{\mu\nu} &= R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} \\ &= \tilde{R}_{\mu\nu} + \delta R_{\mu\nu} - \frac{1}{2} \left(\tilde{R} + \delta R \right) (\tilde{g}_{\mu\nu} + h_{\mu\nu}) \\ &= \tilde{R}_{\mu\nu} - \frac{1}{2} \tilde{R} \tilde{g}_{\mu\nu} + \delta R_{\mu\nu} - \frac{1}{2} \tilde{R} h_{\mu\nu} - \frac{1}{2} \delta R \tilde{g}_{\mu\nu} + O(h^2) \\ &\simeq \tilde{G}_{\mu\nu} + \delta G_{\mu\nu}.\end{aligned}\tag{B.20}$$

We can then inject the expression derived earlier to write (to first order)

$$\begin{aligned}2 \delta G_{\mu\nu} &= 2 \delta R_{\mu\nu} - \delta R \tilde{g}_{\mu\nu} - \tilde{R} h_{\mu\nu} \\ &= \left(\tilde{\nabla}_\alpha \tilde{\nabla}_\nu h_\mu{}^\alpha + \tilde{\nabla}_\alpha \tilde{\nabla}_\mu h_\nu{}^\alpha - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h_{\mu\nu} - \tilde{\nabla}_\nu \tilde{\nabla}_\mu h \right) \\ &\quad - \left(\tilde{\nabla}_\alpha \tilde{\nabla}_\beta h^{\beta\alpha} - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h \right) - \tilde{R} h_{\mu\nu}.\end{aligned}\tag{B.21}$$

To simplify this expression, we will need to perform a few changes and hypotheses. First, it will be useful later to commute the derivatives of the first two terms. The commutation is however not zero, contrarily to usual derivatives. With the Levi-Civita connection, we have that [17, 87]

$$[\tilde{\nabla}_\mu, \tilde{\nabla}_\nu] X^\alpha = \tilde{R}^\alpha{}_{\beta\mu\nu} X^\beta,\tag{B.22}$$

$$[\tilde{\nabla}_\mu, \tilde{\nabla}_\nu] X_\alpha = -\tilde{R}^\beta{}_{\alpha\mu\nu} X_\beta,\tag{B.23}$$

with $[X, Y]$ the commutator, *i.e.* $[X, Y] = XY - YX$. Just as for the connection, we can extend it trivially to several indices, such that

$$[\tilde{\nabla}_\mu, \tilde{\nabla}_\nu] h_\beta{}^\alpha = \tilde{R}^\alpha{}_{\lambda\mu\nu} h_\beta{}^\lambda - \tilde{R}^\lambda{}_{\beta\mu\nu} h_\lambda{}^\alpha.\tag{B.24}$$

In the special case of a repeated index, we get

$$[\tilde{\nabla}_\alpha, \tilde{\nabla}_\nu] h_\mu{}^\alpha = \tilde{R}^\alpha{}_{\lambda\alpha\nu} h_\mu{}^\lambda - \tilde{R}^\lambda{}_{\mu\alpha\nu} h_\lambda{}^\alpha = \tilde{R}_{\lambda\nu} h_\mu{}^\lambda - \tilde{R}^\lambda{}_{\mu\alpha\nu} h_\lambda{}^\alpha.\tag{B.25}$$

If we apply this to the first two terms of Eq. (B.21), we obtain

$$\begin{aligned} \tilde{\nabla}_\alpha \tilde{\nabla}_\nu h_\mu^\alpha + \tilde{\nabla}_\alpha \tilde{\nabla}_\mu h_\nu^\alpha &= \tilde{\nabla}_\nu \tilde{\nabla}_\alpha h_\mu^\alpha + \tilde{R}_{\lambda\nu} h_\mu^\lambda - \tilde{R}^\lambda_{\mu\alpha\nu} h_\lambda^\alpha \\ &+ \tilde{\nabla}_\mu \tilde{\nabla}_\alpha h_\nu^\alpha + \tilde{R}_{\lambda\mu} h_\nu^\lambda - \tilde{R}^\lambda_{\nu\alpha\mu} h_\lambda^\alpha. \end{aligned} \quad (\text{B.26})$$

We can use two properties of the Riemann tensor to simplify this. These properties are the anti-symmetry over the last two indices and the Bianchi identity. We thus have respectively

$$\tilde{R}^\lambda_{\mu\alpha\nu} = -\tilde{R}^\lambda_{\nu\alpha\mu}, \quad (\text{B.27})$$

$$\tilde{R}^\lambda_{\mu\alpha\nu} + \tilde{R}^\lambda_{\alpha\nu\mu} + \tilde{R}^\lambda_{\nu\mu\alpha} = 0. \quad (\text{B.28})$$

We can thus simplify the above expression using both properties, which yields

$$\tilde{R}^\lambda_{\nu\alpha\mu} + \tilde{R}^\lambda_{\mu\alpha\nu} = \tilde{R}^\lambda_{\nu\alpha\mu} - \tilde{R}^\lambda_{\mu\nu\alpha} = -2\tilde{R}^\lambda_{\mu\nu\alpha} - \tilde{R}^\lambda_{\alpha\mu\nu}, \quad (\text{B.29})$$

where we have injected $\tilde{R}^\lambda_{\nu\alpha\mu}$ from the Bianchi identity (Eq. (B.28)). As a result, we can rewrite a couple of terms in Eq. (B.26):

$$-\left(\tilde{R}^\lambda_{\nu\alpha\mu} + \tilde{R}^\lambda_{\mu\alpha\nu}\right) h_\lambda^\alpha = 2\tilde{R}^\lambda_{\mu\nu\alpha} h_\lambda^\alpha + \tilde{R}^\lambda_{\alpha\mu\nu} h_\lambda^\alpha = 2\tilde{R}^\lambda_{\mu\nu\alpha} h_\lambda^\alpha, \quad (\text{B.30})$$

where the second term vanishes because we multiply a symmetric tensor (h) with the anti-symmetric part of another, namely \tilde{R} for which $\tilde{R}_{\lambda\alpha\mu\nu} = -\tilde{R}_{\alpha\lambda\mu\nu}$.

As mentioned, we can also assume that the background metric $\tilde{g}_{\mu\nu}$ is a solution Einstein's equations in the vacuum, so that $\tilde{R} = 0$ and $\tilde{R}_{\mu\nu} = 0$ (solution to $\tilde{G}_{\mu\nu} = 0$). This means that $\tilde{T}_{\mu\nu} = 0$. This is true as long as we are far from any massive body. Thus, Eq. (B.21) becomes

$$\begin{aligned} 2\delta G_{\mu\nu} &= \tilde{\nabla}_\nu \tilde{\nabla}_\alpha h_\mu^\alpha + \tilde{\nabla}_\mu \tilde{\nabla}_\alpha h_\nu^\alpha + 2\tilde{R}^\lambda_{\mu\nu\alpha} h_\lambda^\alpha - \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h_{\mu\nu} - \tilde{\nabla}_\nu \tilde{\nabla}_\mu h \\ &- \tilde{g}_{\mu\nu} \tilde{\nabla}_\alpha \tilde{\nabla}_\beta h^{\beta\alpha} + \tilde{g}_{\mu\nu} \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha h. \end{aligned} \quad (\text{B.31})$$

To simplify this expression further, we will reverse the trace of $h_{\mu\nu}$ by making a change of function. This change of function reads

$$\bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{h}{2} \tilde{g}_{\mu\nu} \quad \bar{h}_\mu{}^\nu = h_\mu{}^\nu - \frac{h}{2} \delta_\mu^\nu \quad \bar{h} = \tilde{g}^{\mu\nu} \bar{h}_{\mu\nu} = -h \quad (\text{B.32})$$

$$\implies h_{\mu\nu} = \bar{h}_{\mu\nu} - \frac{\bar{h}}{2} \tilde{g}_{\mu\nu} \quad h_\mu{}^\nu = \bar{h}_\mu{}^\nu - \frac{\bar{h}}{2} \delta_\mu^\nu \quad (\text{B.33})$$

If we perform this change of function, we get

$$\begin{aligned} 2\delta G_{\mu\nu} &= \tilde{\nabla}_\nu \tilde{\nabla}_\alpha \bar{h}_\mu^\alpha - \tilde{\nabla}_\nu \tilde{\nabla}_\mu \frac{\bar{h}}{2} + \tilde{\nabla}_\mu \tilde{\nabla}_\alpha \bar{h}_\nu^\alpha - \tilde{\nabla}_\mu \tilde{\nabla}_\nu \frac{\bar{h}}{2} + 2\tilde{R}^\lambda_{\mu\nu\alpha} \bar{h}_\lambda^\alpha - 2\tilde{R}^\lambda_{\mu\nu\lambda} \frac{\bar{h}}{2} \\ &- \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha \bar{h}_{\mu\nu} + \tilde{g}_{\mu\nu} \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha \frac{\bar{h}}{2} + \tilde{\nabla}_\nu \tilde{\nabla}_\mu \bar{h} - \tilde{g}_{\mu\nu} \tilde{\nabla}_\alpha \tilde{\nabla}_\beta \bar{h}^{\beta\alpha} \\ &+ \tilde{g}_{\mu\nu} \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha \frac{\bar{h}}{2} - \tilde{g}_{\mu\nu} \tilde{\nabla}_\alpha \tilde{\nabla}^\alpha \bar{h}. \end{aligned} \quad (\text{B.34})$$

We can see that a few terms cancel. We can also note that $\tilde{R}^\lambda{}_{\mu\nu\lambda} = -\tilde{R}^\lambda{}_{\mu\lambda\nu} = \tilde{R}_{\mu\nu} = 0$ by assumption, and that $\tilde{\nabla}_\nu\tilde{\nabla}_\mu\bar{h} = \tilde{\nabla}_\mu\tilde{\nabla}_\nu\bar{h}$, since \bar{h} is a scalar and the connection is symmetric. To be more precise, we have

$$\tilde{\nabla}_\mu\tilde{\nabla}_\nu\bar{h} = \tilde{\nabla}_\mu\partial_\nu\bar{h} = \partial_\mu\partial_\nu\bar{h} - \Gamma^\lambda{}_{\mu\nu}\partial_\lambda\bar{h} = \partial_\nu\partial_\mu\bar{h} - \Gamma^\lambda{}_{\nu\mu}\partial_\lambda\bar{h} = \tilde{\nabla}_\nu\tilde{\nabla}_\mu\bar{h} \quad (\text{B.35})$$

The tensor thus simplifies to

$$2\delta G_{\mu\nu} = \tilde{\nabla}_\nu\tilde{\nabla}_\alpha\bar{h}_\mu{}^\alpha + \tilde{\nabla}_\mu\tilde{\nabla}_\alpha\bar{h}_\nu{}^\alpha + 2\tilde{R}^\lambda{}_{\mu\nu\alpha}\bar{h}_\lambda{}^\alpha - \tilde{\nabla}_\alpha\tilde{\nabla}^\alpha\bar{h}_{\mu\nu} - \tilde{g}_{\mu\nu}\tilde{\nabla}_\alpha\tilde{\nabla}_\beta\bar{h}^{\beta\alpha}. \quad (\text{B.36})$$

We will do the last simplification that will cancel most of the remaining terms. We choose to use the De Donder gauge

$$\tilde{\nabla}_\alpha\bar{h}^{\alpha\beta} = 0, \quad (\text{B.37})$$

such that the tensor finally becomes

$$2\delta G_{\mu\nu} = -\tilde{\nabla}_\alpha\tilde{\nabla}^\alpha\bar{h}_{\mu\nu} + 2\tilde{R}^\lambda{}_{\mu\nu\alpha}\bar{h}_\lambda{}^\alpha. \quad (\text{B.38})$$

We can see that this expression is much simpler than Eq. (B.34).

Hence, under the assumption that second or higher-order terms are negligible, with $\tilde{g}_{\mu\nu}$ that is a solution to the vacuum Einstein's equations (*i.e.* with $\tilde{T}_{\mu\nu} = 0$) and using the De Donder gauge, the equations become

$$\tilde{G}_{\mu\nu} + \delta G_{\mu\nu} = 8\pi G \left(\tilde{T}_{\mu\nu} + \delta T_{\mu\nu} \right) \quad (\text{B.39})$$

$$\implies \delta G_{\mu\nu} = \frac{1}{2} \left(-\tilde{\nabla}_\alpha\tilde{\nabla}^\alpha\bar{h}_{\mu\nu} + 2\tilde{R}^\lambda{}_{\mu\nu\alpha}\bar{h}_\lambda{}^\alpha \right) = 8\pi G \delta T_{\mu\nu}, \quad (\text{B.40})$$

where $\delta T_{\mu\nu}$ is related to the source of the gravitational wave. If we assume that we are outside the source, we can take $\delta T_{\mu\nu} = 0$. We then obtain the *linearised Einstein's equations in vacuum* and the full system including the gauge condition reads

$$\begin{cases} -\tilde{\nabla}_\alpha\tilde{\nabla}^\alpha\bar{h}_{\mu\nu} + 2\tilde{R}^\lambda{}_{\mu\nu\alpha}\bar{h}_\lambda{}^\alpha = 0, \\ \tilde{\nabla}_\alpha\bar{h}^{\alpha\beta} = 0. \end{cases} \quad (\text{B.41})$$

It can also be found written as [18]

$$\begin{cases} \tilde{\nabla}_\alpha\tilde{\nabla}^\alpha\bar{h}_{\mu\nu} + 2\tilde{R}^\lambda{}_{\mu\alpha\nu}\bar{h}_\lambda{}^\alpha = 0, \\ \tilde{\nabla}_\alpha\bar{h}^{\alpha\beta} = 0. \end{cases} \quad (\text{B.42})$$

This is obtained by simply changing the order of the two last indices of \tilde{R} , which leads to the apparition of a minus sign.

Appendix C

Rewriting the wave equation

To rewrite the wave equation, one starts from

$$\nabla_\alpha \nabla^\alpha h_{\mu\nu} = \nabla_\alpha \nabla^\alpha (\Phi e_{\mu\nu}) = 0. \quad (\text{C.1})$$

Assuming the polarisation does not change, one can then write

$$\nabla_\alpha \nabla^\alpha \Phi = \nabla_\alpha \partial^\alpha \Phi = 0, \quad (\text{C.2})$$

where we used the fact that the covariant derivative is the partial derivative when applied to a scalar. Considering any vector V^α , one can show that [87]

$$\nabla_\alpha V^\alpha = \frac{1}{\sqrt{-g}} \partial_\alpha (\sqrt{-g} V^\alpha), \quad (\text{C.3})$$

where g is the determinant of the metric. Note that one needs a minus sign, since the determinant is negative (see Eq. the Minkowski metric). In our case, $V^\alpha = \partial^\alpha \Phi$. Using also the fact that $\partial^\alpha = g^{\alpha\lambda} \partial_\lambda$, the wave equation then becomes

$$\nabla_\alpha \nabla^\alpha \Phi = \frac{1}{\sqrt{-g}} \partial_\alpha (\sqrt{-g} g^{\alpha\lambda} \partial_\lambda \Phi) = 0. \quad (\text{C.4})$$

Bibliography

- [1] R. A. Hulse and J. H. Taylor. ‘Discovery of a pulsar in a binary system.’ In: *The Astrophysical Journal* 195 (1975), pp. L51–L53.
- [2] J. H. Taylor and J. M. Weisberg. ‘A new test of general relativity - Gravitational radiation and the binary pulsar PSR 1913+16’. In: *The Astrophysical Journal* 253 (1982), pp. 908–920.
- [3] B. P. Abbott et al. ‘Observation of Gravitational Waves from a Binary Black Hole Merger’. In: *Phys. Rev. Lett.* 116 (2016), p. 061102.
- [4] The LIGO Scientific Collaboration, The Virgo Collaboration and The KAGRA Collaboration. *GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run*. 2021. URL: <https://arxiv.org/abs/2111.03606>.
- [5] The LIGO Scientific Collaboration et al. ‘Advanced LIGO’. In: *Classical and Quantum Gravity* 32.7 (2015), p. 074001.
- [6] B. P. Abbott et al. ‘LIGO: the Laser Interferometer Gravitational-Wave Observatory’. In: *Reports on Progress in Physics* 72.7 (2009), p. 076901.
- [7] T. Accadia et al. ‘Status of the Virgo project’. In: *Classical and Quantum Gravity* 28.11 (2011), p. 114002.
- [8] F. Acernese et al. ‘Advanced Virgo: a second-generation interferometric gravitational wave detector’. In: *Classical and Quantum Gravity* 32.2 (2014), p. 024001.
- [9] Y. Aso et al. ‘Interferometer design of the KAGRA gravitational wave detector’. In: *Phys. Rev. D* 88 (2013), p. 043007.
- [10] K. L. Dooley et al. ‘GEO 600 and the GEO-HF upgrade program: successes and challenges’. In: *Classical and Quantum Gravity* 33.7 (2016), p. 075009.
- [11] D. Walsh et al. ‘0957 + 561 A, B: twin quasistellar objects or gravitational lens?’ In: *Nature* 279.5712 (1979), pp. 381–384.
- [12] R. Abbott et al. ‘Search for Lensing Signatures in the Gravitational-Wave Observations from the First Half of LIGO–Virgo’s Third Observing Run’. In: *The Astrophysical Journal* 923.1 (2021), p. 14.
- [13] M. Punturo et al. ‘The Einstein Telescope: a third-generation gravitational wave observatory’. In: *Classical and Quantum Gravity* 27.19 (2010), p. 194002.

- [14] D. Reitze et al. ‘Cosmic Explorer: The U.S. Contribution to Gravitational-Wave Astronomy beyond LIGO’. In: *Bulletin of the AAS* 51.7 (2019), p. 35.
- [15] M. Maggiore et al. ‘Science case for the Einstein telescope’. In: *Journal of Cosmology and Astroparticle Physics* 2020.3 (2020), p. 50.
- [16] E. Cuoco et al. ‘Enhancing gravitational-wave science with machine learning’. In: *Machine Learning: Science and Technology* 2.1 (2020), p. 011002.
- [17] S. M. Carroll. *Lecture Notes on General Relativity*. 1997. URL: <https://arxiv.org/abs/gr-qc/9712019>.
- [18] C. W. Misner, K. S. Thorne and J. A. Wheeler. *Gravitation*. Princeton University Press, 2017.
- [19] D. Baumann. *Cosmology*. (Lecture notes). URL: <http://cosmology.amsterdam/education/cosmology/> (visited on 03/05/2022).
- [20] D. Tong. *Lectures on Cosmology*. 2019. URL: <http://www.damtp.cam.ac.uk/user/tong/cosmo.html> (visited on 03/05/2022).
- [21] A. G. Riess et al. ‘Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant’. In: *The Astronomical Journal* 116.3 (1998), pp. 1009–1038.
- [22] M. Meneghetti. *Introduction to Gravitational Lensing. With Python Examples*. Springer, Cham, 2021.
- [23] A.O. Petters, H. Levine and J. Wambsganss. *Singularity Theory and Gravitational Lensing*. Birkhäuser, Boston, MA, 2001.
- [24] P. Schneider, J. Ehlers and E. E. Falco. *Gravitational lenses*. Springer, Berlin, Heidelberg, 1992.
- [25] A. Le Tiec and J. Novak. ‘Theory of Gravitational Waves’. In: *An Overview of Gravitational Waves: Theory*. 2017, pp. 1–41.
- [26] V. Perlick. *Gravitational Waves*. (Lecture notes). 2017. URL: <https://www.zarm.uni-bremen.de/en/research/space-science/gravitational-theory/team-members/personal-websites/volker-perlick.html> (visited on 16/05/2022).
- [27] M. Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press, 2008.
- [28] J.D.E. Creighton and W.G. Anderson. *Gravitational-Wave Physics and Astronomy: An Introduction to Theory, Experiment and Data Analysis*. John Wiley & Sons, Ltd, 2011.
- [29] O. Bulashenko and H. Ubach. *Lensing of gravitational waves: universal signatures in the beating pattern*. 2021. URL: <https://arxiv.org/abs/2112.10773>.
- [30] R. Takahashi and T. Nakamura. ‘Wave Effects in the Gravitational Lensing of Gravitational Waves from Chirping Binaries’. In: *The Astrophysical Journal* 595.2 (2003), pp. 1039–1051.

- [31] E. T. Copson. *Asymptotic Expansions*. Cambridge Tracts in Mathematics. Cambridge University Press, 1965.
- [32] M. Born and E. Wolf. *Principles of Optics: 60th Anniversary Edition*. 7th ed. Cambridge University Press, 2019.
- [33] A. S. Davydov. *Quantum mechanics*. International series of monographs in natural philosophy, 1. Oxford: Pergamon, 1965.
- [34] P. C. Peters. ‘Index of refraction for scalar, electromagnetic, and gravitational waves in weak gravitational fields’. In: *Physical Review D* 9.8 (1974), pp. 2207–2218.
- [35] S. Deguchi and W. D. Watson. ‘Diffraction in Gravitational Lensing for Compact Objects of Low Mass’. In: *The Astrophysical Journal* 307 (1986), p. 30.
- [36] *NIST Digital Library of Mathematical Functions*. <http://dlmf.nist.gov/>, Release 1.1.5 of 2022-03-15. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds. URL: <http://dlmf.nist.gov/>.
- [37] S. Deguchi and W. D. Watson. ‘Wave effects in gravitational lensing of electromagnetic radiation’. In: *Physical Review D* 34.6 (1986), pp. 1708–1718.
- [38] K. Liao et al. ‘The Wave Nature of Continuous Gravitational Waves from Microlensing’. In: *The Astrophysical Journal* 875.2 (2019), p. 139.
- [39] G. Louppe. *Lectures on deep learning*. URL: <https://github.com/glouppe/info8010-deep-learning> (visited on 04/05/2022).
- [40] T. Hastie, R. Tibshirani and J. Friedman. *The Elements of Statistical Learning. Data mining, Inference, and Prediction*. 2nd ed. Springer, New York, NY, 2009.
- [41] A. Zhang et al. ‘Dive into Deep Learning’. In: *arXiv preprint arXiv:2106.11342* (2021). URL: <https://d2l.ai/index.html>.
- [42] M. Belkin et al. *Reconciling modern machine learning practice and the bias-variance trade-off*. 2018. URL: <https://arxiv.org/abs/1812.11118>.
- [43] Y. LeCun et al. ‘Deep learning’. In: *Nature* 521 (2015), pp. 436–444.
- [44] A.G. Baydin et al. ‘Automatic Differentiation in Machine Learning: a Survey’. In: *Journal of machine learning research* 18.153 (2018), pp. 1–43.
- [45] A. Nitz et al. *gwastro/pycbc: v2.0.2 release of PyCBC*. Version v2.0.2. 2022. URL: <https://doi.org/10.5281/zenodo.6324278>.
- [46] LIGO Scientific Collaboration. *LIGO Algorithm Library - LALSuite*. free software (GPL). 2018.
- [47] A. Bohé et al. ‘Improved effective-one-body model of spinning, nonprecessing binary black holes for the era of gravitational-wave astrophysics with advanced detectors’. In: *Phys. Rev. D* 95 (2017), p. 044028.
- [48] J. Janquart. ‘Gravitational waves signal analysis: Matched filtering, typical analyses and beyond’. MA thesis. 2020. URL: <https://matheo.uliege.be/handle/2268.2/9211?locale=en> (visited on 03/05/2022).

- [49] C. Cutler and É. E. Flanagan. ‘Gravitational waves from merging compact binaries: How accurately can one extract the binary’s parameters from the inspiral waveform?’ In: *Phys. Rev. D* 49 (1994), pp. 2658–2697.
- [50] B. J. Owen and B. S. Sathyaprakash. ‘Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement’. In: *Phys. Rev. D* 60 (1999), p. 022002.
- [51] B.J. Owen. ‘Search templates for gravitational waves from inspiraling binaries: Choice of template spacing’. In: *Phys. Rev. D* 53 (1996), pp. 6749–6761.
- [52] B. Allen et al. ‘FINDCHIRP: An algorithm for detection of gravitational waves from inspiraling compact binaries’. In: *Phys. Rev. D* 85 (2012), p. 122006.
- [53] P. G. Krastev et al. ‘Detection and parameter estimation of gravitational waves from binary neutron-star mergers in real LIGO data using deep learning’. In: *Physics Letters B* 815 (2021), p. 136161.
- [54] K. Cannon et al. ‘Toward early-warning detection of gravitational waves from compact binary coalescence’. In: *The Astrophysical Journal* 748.2 (2012), p. 136.
- [55] B. P. Abbott et al. ‘Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A’. In: *The Astrophysical Journal* 848.2 (2017), p. L13.
- [56] B. P. Abbott et al. ‘GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral’. In: *Phys. Rev. Lett.* 119 (2017), p. 161101.
- [57] J. Yan et al. ‘Generalized approach to matched filtering using neural networks’. In: *Physical Review D* 105.4 (2022), p. 043006.
- [58] H. Gabbard et al. ‘Matching Matched Filtering with Deep Networks for Gravitational-Wave Astronomy’. In: *Phys. Rev. Lett.* 120 (2018), p. 141103.
- [59] D. George and E.A. Huerta. ‘Deep Learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data’. In: *Physics Letters B* 778 (2018), pp. 64–70.
- [60] J. D. Álvares et al. ‘Exploring gravitational-wave detection and parameter inference using deep learning methods’. In: *Classical and Quantum Gravity* 38.15 (2021), p. 155010.
- [61] T. Marianer et al. ‘A semisupervised machine learning search for never-seen gravitational-wave sources’. In: *Monthly Notices of the Royal Astronomical Society* 500.4 (2020), pp. 5408–5419.
- [62] S. Goyal et al. ‘Rapid Identification of Strongly Lensed Gravitational-Wave Events with Machine Learning’. In: *Phys. Rev. D* 104 (2021), p. 124057.
- [63] Kim K. et al. ‘Identification of Lensed Gravitational Waves with Deep Learning’. In: *The Astrophysical Journal* 915.2 (2021), p. 119.
- [64] K. Haris et al. *Identifying strongly lensed gravitational wave signals from binary black hole mergers*. 2018. URL: <https://arxiv.org/abs/1807.07062>.

- [65] G. Huang et al. ‘Densely Connected Convolutional Networks’. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017, pp. 4700–4708.
- [66] R. Abbott et al. ‘GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo during the First Half of the Third Observing Run’. In: *Phys. Rev. X* 11 (2021), p. 021053.
- [67] B.P. Abbott et al. ‘Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA’. In: *Living Reviews in Relativity* 23 (2020), p. 3.
- [68] D. P. Kingma and J. Ba. *Adam: A Method for Stochastic Optimization*. 2014. URL: <https://arxiv.org/abs/1412.6980>.
- [69] M. Çalışkan et al. *Lensing or luck? False alarm probabilities for gravitational lensing of gravitational waves*. 2022. URL: <https://arxiv.org/abs/2201.04619>.
- [70] O. A. Hannuksela et al. ‘Localizing merging black holes with sub-arcsecond precision using gravitational-wave lensing’. In: *Monthly Notices of the Royal Astronomical Society* 498.3 (2020), pp. 3395–3402.
- [71] X.-L. Fan et al. ‘Speed of Gravitational Waves from Strongly Lensed Gravitational Waves and Electromagnetic Signals’. In: *Phys. Rev. Lett.* 118 (2017), p. 091102.
- [72] K. Liao et al. ‘Precision cosmology from future lensed gravitational wave and electromagnetic signals’. In: *Nature Communications* 8 (2017), p. 1148.
- [73] J. M. Ezquiaga and M. Zumalacárregui. ‘Gravitational wave lensing beyond general relativity: Birefringence, echoes, and shadows’. In: *Phys. Rev. D* 102 (2020), p. 124048.
- [74] F. Xu et al. ‘Please Repeat: Strong Lensing of Gravitational Waves as a Probe of Compact Binary and Galaxy Populations’. In: *The Astrophysical Journal* 929.1 (2022), p. 9.
- [75] L. Dai et al. ‘Detecting lensing-induced diffraction in astrophysical gravitational waves’. In: *Phys. Rev. D* 98 (2018), p. 104029.
- [76] M. Oguri. ‘Effect of gravitational lensing on the distribution of gravitational waves from distant binary black hole mergers’. In: *Monthly Notices of the Royal Astronomical Society* 480.3 (2018), pp. 3842–3855.
- [77] B. P. Abbott et al. ‘The basic physics of the binary black hole merger GW150914’. In: *Annalen der Physik* 529.1-2 (2016), p. 1600209.
- [78] L. Dai et al. ‘Effect of lensing magnification on the apparent distribution of black hole mergers’. In: *Phys. Rev. D* 95 (2017), p. 044011.
- [79] M. Fishbach et al. ‘Are LIGO’s Black Holes Made from Smaller Black Holes?’ In: *The Astrophysical Journal* 840.2 (2017), p. L24.
- [80] T. Broadhurst et al. *Reinterpreting Low Frequency LIGO/Virgo Events as Magnified Stellar-Mass Black Holes at Cosmological Distances*. 2018. URL: <https://arxiv.org/abs/1802.05273>.

- [81] G. P. Smith et al. ‘What if LIGO’s gravitational wave detections are strongly lensed by massive galaxy clusters?’ In: *Monthly Notices of the Royal Astronomical Society* 475.3 (2018), pp. 3823–3828.
- [82] K. K. Y. Ng et al. ‘Precise LIGO lensing rate predictions for binary black holes’. In: *Phys. Rev. D* 97 (2018), p. 023012.
- [83] S.-S. Li et al. ‘Gravitational lensing of gravitational waves: a statistical perspective’. In: *Monthly Notices of the Royal Astronomical Society* 476.2 (2018), pp. 2220–2229.
- [84] M. B. Schäfer et al. ‘Training strategies for deep learning gravitational-wave searches’. In: *Phys. Rev. D* 105 (2022), p. 043002.
- [85] A. Vaswani et al. ‘Attention is All you Need’. In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017, pp. 5998–6008.
- [86] A. Dosovitskiy et al. *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*. 2020. URL: <https://arxiv.org/abs/2010.11929>.
- [87] M. Blau. *Lecture Notes on General Relativity*. 2022. URL: <http://www.blau.itp.unibe.ch/GRlecturenotes.html> (visited on 16/05/2022).