



Universitat
de les Illes Balears

Towards a Reformulation of the Hough Method for Continuous Gravitational Wave Searches

Rodrigo Tenorio Márquez

Master's Thesis

Master's degree in Advanced Physics and Applied Mathematics
at the

UNIVERSITAT DE LES ILLES BALEARS

Academic year 2018-2019

Date: June, 2019

UIB Master's Thesis Supervisor: Dra. Alicia M. Sintes Olives

Abstract

The recent detections of gravitational waves performed by the Advanced LIGO and Advanced Virgo interferometers have established an entirely new channel to probe information from the Universe. In a complementary fashion to electromagnetic waves, gravitational waves are perturbations of the spacetime, traveling through the Universe.

So far, the detected gravitational waves were produced as a result of the coalescence of compact objects, such as black holes or neutron stars. They emit a highly concentrated burst of energy, following a well-understood propagation mechanism. However, there is more.

It is theoretically described that some compact objects are able to sustain imperfections in its shape for a long period of time. The most common example is given by a neutron star, the crust of which may present a small *bump* which corresponds to a deviation from a perfectly symmetric object. Due to the persistence in time of such physical mechanisms, the emitted signal presents a remarkable property with respect with the previous situation: It may last for long periods of time, compared to the extremely short duration of the coalescence of compact binaries; thus, we label these signals as *continuous waves*.

The detection of such continuous signals will open a new chapter in gravitational wave astronomy, extending the field from transient signals to the observation of persistent sources, which can be monitored with ever increasing precision, over very long periods of time, as is common in traditional electromagnetic astronomy. Due to the modulation of the signal resulting from the motion of the Earth, it will be possible to accurately measure both gravitational wave polarizations, and the accurate tracking of the gravitational wave phase over many years will yield information about the equation of state and possible transient events like seismic activity in the neutron star crust.

The work presented here further develops the capabilities of the *SkyHough* pipeline, which is one of the main tools used by the LIGO and Virgo Collaboration to search for continuous wave signals. This search algorithm implements the Hough transform, which is a type of pattern recognition algorithm, first developed to recognize particle tracks in bubble chambers. One of the main advantages of this method has been its robustness against noise.

As opposed to current searches for transients, like comparable mass compact objects, searches for continuous wave signals are limited by computational resources, and carrying out optimally sensitive searches is computationally prohibitive. All-sky searches are therefore typically carried out semicoherently: a stream of data is first split into different time segments, then a suitable quantity measuring significance is assigned to each segment. Finally, the results for each segment are combined into a single quantity, which can be used to state the presence of a signal within the analyzed dataset.

Noise artifacts populate every step of the statistical procedure. Even though one can attempt to construct an analytical derivation of the noise statistical properties at each step, it is more reliable to work with

numerical approaches, as they take into account the effects measured by the data analysis procedure. A good representation of the noise distribution is important to state the significance of a certain result.

The main contribution of this work is the development of a new formulation which combines data from different detectors into a single analysis. This new method relies on a more accurate estimation of the noise, which replaces the assumption of some prescribed underlying noise distribution, as has been used in previous versions of the *SkyHough* pipeline. This noise estimation is based on the introduction of an efficient random sampling procedure in the large parameter space of signals. Working with this new accurate estimate of the background noise distribution, one is able to increase the sensitivity of the searches, as a better understanding of the background behavior translates to a better identification of signals in terms of its statistical significance.

In addition, I have started to explore further improvements to the *SkyHough* pipeline: the use of the universal statistics approach of [26], the use of artificial neural networks for candidate classification, and the utilization of a method developed by [19] to deal with spectral leakage. These investigations are not yet concluded, but allow us to explore the directions in which data analysis pipelines for continuous wave searches may evolve in the near future.

Together with the continuously increasing sensitivity of our detectors, it is expected that these improvements will contribute to keep the *SkyHough* pipeline at the forefront of data analysis of continuous waves, pointing us towards the correct way to a first detection.

The document is structured as follows: Chapter 1 introduces the basic theory on gravitational waves and interferometric detectors; chapter 2 introduces neutron stars, discussing their relevance within the field of continuous gravitational waves; chapter 3 describes the *SkyHough* pipeline for the detection of continuous gravitational waves; chapter 4 introduces the S6 Mock Data Challenge, which is the data set that we used to test our improvements against previous implementations of the *SkyHough* pipeline; chapter 5 discusses the main results, implementing the sampling procedure to properly take into account the contribution of noise fluctuations into the data analysis pipeline; chapter 6 describes some minor developments, still on an experimental stage, which could yield relevant contributions to the data analysis procedure; chapter 7 concludes the work, summarizing the main results.

Acknowledgements

I would like to thank my thesis advisor, Dra. Alicia Sintes, for introducing me to research and for her guidance and help during this year.

I would like to thank Dr. Sascha Husa, for all his invaluable lessons on the proper way to do research.

I would also like to thank the continuous wave subdivision of the LIGO@UIB group, Miquel and Pep, for many fruitful discussions about the analysis of LIGO data. I also thank Rafel, for being a constant support during the development of this work.

I am grateful to the population of the CBC-cave for their hard-working spirit, continuously learning one from another to deliver the best science they can.

Finally, I would like to express my sincere gratitude to my family for always supporting my choices, even when they were not as obvious as they would expect.

Contents

I	Introduction	9
1	Introduction to Gravitational Waves	11
1.1	General Relativity: Linearized Theory	11
1.2	Emission of Gravitational Waves	12
1.3	Detection of Gravitational Waves	14
1.3.1	Power Spectral Density of a Detector	15
1.3.2	General Noise Contributions	17
1.3.3	Persistent Spectral Disturbances: Lines	18
2	Neutron Stars as Gravitational Wave Sources	19
2.1	Physics of a Neutron Star	19
2.1.1	Emission Mechanisms	20
2.1.2	Population of Neutron Stars	21
2.2	Signal from an Isolated Neutron Star	23
2.2.1	Phase and Frequency: From the Source to the Detector	23
2.2.2	Amplitude of the Gravitational Wave	24
2.3	Data Analysis for a Continuous Gravitational Waves	25
2.3.1	Coherent Methods	26
2.3.2	Semicoherent methods	26
2.4	State of the Art on Continuous Wave Searches	27
2.4.1	All Sky Search	27
2.4.2	Targeted Search	27
2.4.3	Narrow Band Search	30
2.4.4	Directed Search	31
2.5	Summary	31
3	The Hough Transform for Continuous Gravitational Waves Searches	33
3.1	Overview of the Standard SkyHough	34
3.1.1	Coherent Step	34
3.1.2	Incoherent Step	35
3.2	Implementation	36
3.2.1	Length of Coherent Segments	37
3.2.2	Discretization of the Parameter Space	38
3.2.3	Spin Down ranges	38
3.2.4	Partial Hough Maps	39
3.3	Statistics	40

3.3.1	Normalized Power distribution	41
3.3.2	Number Count Distribution	42
3.3.3	Optimal Choice of thresholds	43
3.4	Post Processing	44
3.4.1	Parameter Space Distance	44
3.4.2	Coincidence and Clustering	45
4	The sixth LIGO scientific run and its Continuous Wave Mock Data Challenge	47
4.1	LIGO S6 Data	47
4.2	The S6 Mock Data Challenge	49
4.2.1	The Mock Data Challenge as a Test Dataset	50
II	Original Results	53
5	A Noise Robust Statistic for the SkyHough Pipeline	55
5.1	The Robust Statistic	55
5.2	The Multi-Interferometer Statistic	57
5.3	Implementation of the Robust Statistic	58
5.3.1	Parameter Space Sampling	58
5.3.2	Estimation of Sample Statistics	59
5.3.3	Multi-Interferometer Number Count	59
5.4	Performance of the Robust Statistic	60
5.4.1	Analysis of Known Lines	60
5.4.2	S6 Mock Data Challenge: Overview	62
5.4.3	Efficiency and Parameter Recovery	65
6	Future Work	69
6.1	Tau Statistic	69
6.1.1	Implementation	71
6.2	Universal Statistic	71
6.3	Post Processing Reformulation	72
6.4	Summary	73
7	Conclusion	75
	Bibliography	77

Part I

Introduction

Introduction to Gravitational Waves

1.1 General Relativity: Linearized Theory

Einstein's general theory of relativity introduces a geometric approach to the description of gravity, unifying the classical conceptions of space and time into a single object, the *spacetime*. As a result, gravity and inertia are sewed together as a pure manifestation of the spacetime geometry.

The actual interaction between matter and geometry is given in terms of the Einstein Field Equations

$$G_{ab} = 8\pi T_{ab}, \quad (1.1)$$

a tensorial equation which involves the Einstein tensor G_{ab} and the stress-energy tensor T_{ab} . The former describes the spacetime geometry, while the latter describes the way in which matter flows within it; in some sense, matter tells spacetime how to curve, and spacetime tells matter how to flow. The equation is written using natural units: The speed of light is set to $c = 1$, as well as the gravitational constant $G = 1$, allowing length, time and mass to be described using the same units.

The Einstein tensor is described in terms of a more fundamental object, the metric tensor g_{ab} , due to which spacetime is endowed with a geometry. However, the functional relation between G_{ab} and g_{ab} is highly non-linear, difficulting the task of finding analytical solutions.

Nevertheless, one may try to develop a linearized version of the theory, working on *perturbations* over a background metric. For instance, we could choose to set a *flat*, empty background, where free particles move following straight lines; such a space would be described by the Minkowski metric η_{ab} and a null stress-energy tensor, $T_{ab} = 0$. Introducing this considerations into (1.1) and expanding up to linear order in the perturbation one ends up with a wave equation

$$\square h_{ab}^{TT} = 0, \quad (1.2)$$

where \square represents the D'Alembertian operator and the label TT denotes the choice of a specific gauge condition. This wave-like perturbations are called Gravitational Waves (GW).

Being a general covariance theory, general relativity allows a free choice of coordinates to describe the ongoing physics. For instance, the aforementioned Transverse Traceless (TT) gauge is well suited to describe traveling waves, since they become pure spacetime deformations perpendicular to the direction of propagation. Choosing a Cartesian frame of reference (t, x, y, z) , the amplitude of a GW traveling




Figure 1.1: Effect of the different polarizations of a gravitational wave traveling across a ring of test masses along the normal direction. They produce perturbations towards perpendicular directions in an *alternated* fashion. Both polarizations differ by a rotation of 45° , hence its names. Source: [78]

towards $+z$ would be given by

$$h_{\mu\nu}^{TT} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (1.3)$$

where $h_{+,\times}$ stand for the two allowed polarizations of a gravitational wave. We refer to [53] for a detailed derivation of this fact. Figure 1.1 illustrates the behavior of both polarizations on a ring of test masses.

The choice of the TT gauge implies a particular frame of reference: The coordinate curves are trajectories of free falling, test masses. That leads to the following result: A pair of test masses initially at rest with respect to each other *before* the transient of a GW will still be at rest even *after* the transient, since their coordinates, given by its own trajectories, will remain *unchanged*. However, if we were to compute the actual *distance* between the same pair of masses, we would realize a variation due to the gravitational wave transient, since the spacetime *metric* is perturbed. Hence, from the TT frame, gravitational waves deform spacetime.

There exist a second option to model the effects of gravitational waves, more suited to the experimental measurements: Regard its effects as classical, tidal forces. This can be done by constructing a coordinate system using rigid rods¹, actually encapsulating the effect of gravitational waves into the displacement between test masses, which are pushed and pulled by a fictitious force. Such a displacement is currently measured by means of laser interferometry, as we will discuss during the following sections.

1.2 Emission of Gravitational Waves

So far, we have described the traveling process of gravitational waves. If we attempt to describe its *emission*, we will inevitably face a major issue, since our linearized theory turns out invalid in the strong regime that dominates such a process.

¹Given a rod with length L , the deformation imprinted by the transit of a GW with frequency ω is given by $\Delta L \sim \left(\frac{\omega}{\omega_0}\right)^2$, where $\omega_0 \sim L^{-1}$ is the fundamental frequency of the rod. Hence, the assumption on rigid rods is valid as long as we consider them to be as short as needed.

In order to circumvent this problem, we will grasp the production of gravitational waves through a direct comparison with electromagnetic theory. Thanks to the similarities between Newton's and Coulomb's descriptions, taking $e^2 \rightarrow -m^2$ will suffice this purpose. This duality among classical field theories can be used to identify further variables and impose further constraints into possible emission mechanisms through the use of conservation laws.

First, as does electrical charge, mass satisfies conservation principle, which forbids any form of *monopolar* radiation. As for the dipole radiations, we could attempt to write analogous to the electric and magnetic dipole radiations: The first one relies on linearly accelerating charges, which translate into accelerating bodies through the previously stated duality; due to global lineal momentum conservation, this emission channel becomes forbidden. As for the second one, it can be described in terms of current loops, which becomes related to the total angular momentum of a system; again, conservation of global angular momentum vetoes this mechanism. Hence, there is no *dipolar* gravitational radiation.

We need to reach quadrupolar radiation order achieve a feasible mechanism of gravitational waves emission. Defining \mathcal{J}_{ab} as the mass quadrupole moment, completely analogous to its electromagnetic version, the metric perturbation in TT gauge for a mass distribution will be given by

$$h_{ab}^{TT}(t, \vec{r}) = \frac{1}{r} \frac{2G}{c^4} \ddot{\mathcal{J}}_{ab}(t - r/c), \quad (1.4)$$

where \vec{r} represents the position from the source and time differentiation is represented by dots.

The rate of energy emitted through quadrupolar radiation can be expressed as the gravitational luminosity \mathcal{L}_{GW} in terms of third temporal derivatives on the quadrupolar mass tensor

$$\mathcal{L}_{GW} = \frac{1}{5} \frac{G}{c^5} \langle \ddot{\mathcal{J}}_{ab} \ddot{\mathcal{J}}^{ab} \rangle. \quad (1.5)$$

Simply stated, any non-axisymmetric mass distribution can potentially emit gravitational waves.

Let us consider particular case of an *isolated, quasi-spherical*, spinning object with mass M and radius R . *Quasi-sphericity* will be achieved by breaking the degeneracy on the three principal moments of inertia I_1, I_2, I_3 , using the equatorial *ellipticity* as an adimensional quantity of such a magnitude

$$\epsilon = \frac{I_1 - I_2}{I_3}, \quad (1.6)$$

being $\epsilon = 0$ the case of an axisymmetric object. Its gravitational luminosity will be given by (1.5)

$$\mathcal{L}_{GW} \sim \frac{G}{c^5} \epsilon^2 I_3^2 \nu^6, \quad (1.7)$$

where ν stands for the spinning frequency of the object. In order to express the order of magnitude in a useful manner, we will refer distances to the Schwarzschild radius $R_s = 2GMc^{-2}$ and rewrite frequencies in terms of the spinning velocity of the body $v = 2\pi R\nu$. All along, we obtain

$$\mathcal{L}_{GW} \sim \frac{c^5}{G} \epsilon^2 \left(\frac{R_s}{R} \right)^2 \left(\frac{v}{c} \right)^6, \quad (1.8)$$

which clearly indicates that the most suitable gravitational wave sources are compact objects (i.e. $R \sim R_s$) rotating at high velocities (i.e. $v \sim c$) with a sustained ellipticity ϵ . Being $c^5/G \sim 10^{52} W$, even small departures from axisymmetry will imply a significant flow of energy.

The related amplitude of the gravitational wave can be derived from (1.4)

$$h_0 \sim 3 \cdot 10^{-25} \left(\frac{\epsilon}{10^{-6}} \right) \left(\frac{I_3}{10^{38} \text{kgm}^2} \right) \left(\frac{\nu}{100 \text{Hz}} \right)^2 \left(\frac{100 \text{pc}}{r} \right), \quad (1.9)$$

where the scales are taken in a compatible fashion with nearby neutron stars in our galaxy. It becomes clear that gravitational waves represent a major trial on experimental physics, due to the extremely small effects over terrestrial objects.


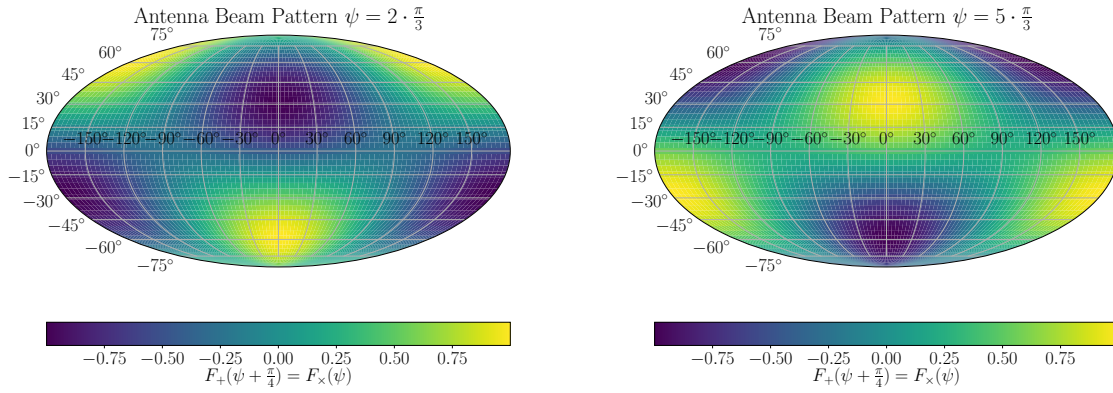


Figure 1.2: Schematic view of an interferometric gravitational wave detector. Source: [63]


 Figure 1.3: Molweide projection of the Antenna Beam Pattern of an interferometric detector, showing the response to gravitational waves coming from sources with two particular polarizations. Values close to ± 1 indicate that the gravitational wave produces a differential motion over the two arms, while null values correspond to a gravitational wave which creates a *consonant* perturbation over both arms.

1.3 Detection of Gravitational Waves

Nowadays, the detection of gravitational waves is performed using interferometric detectors. The idea is to set up the mirrors that conform the interferometer as *free falling* masses; that way, a laser beam can be used to measure spacetime perturbations following the spirit of previous section.

The actual measurement is performed by taking into account the *alternated* perturbation of mirrors located at both interferometers. As discussed in figure 1.1, gravitational waves will disturb the position of free falling masses in a alternated fashion, according to its two polarizations $h_{+,\times}$. Taking the separation between the beam splitter and the end mirrors to be L , the perturbation due to a gravitational wave can be expressed in terms of the length variation ΔL as the *strain*

$$h(t) \equiv \frac{\Delta L(t)}{L} = F_+(\vec{n}, \psi; t)h_+(t) + F_\times(\vec{n}, \psi; t)h_\times(t). \quad (1.10)$$

$F_{+,\times}$ are the antenna beam pattern functions, which represent the response of the detector to a particular wave polarization coming from the sky position \vec{n} with the main axis of its $+$ polarization twisted by



Figure 1.4: Interferometric gravitational wave detectors. **Left and Center:** LIGO Livingston and LIGO Hanford observatories, located in the United States of America. **Right:** Virgo observatory, located in Italy. Source: [50, 85].

an angle of ψ . Figure 1.3 shows the detector response to a certain polarization as a function of the sky position. The exact form of F_+ and F_\times can be found in [41].

By principle, interferometric detectors are omnidirectional, as opposed to telescopes, which must be pointed towards a certain sky position: They act alike *microphones*. This is advantageous, since *almost* every single sky position is covered at once, being the exception the points in which the antenna pattern response is low,

Although it is an omnidirectional instrument, sky locations can be easily determined using the suitable strategy: For the case of a *continuous* signal (i.e. a persistent signal in time), the sky position can be tracked down using the Doppler effect introduced by the Earth's movement and the anisotropic detector response; as for the case of a transient signal, multiple detectors can be used to infer the position of the source.

Nowadays, there are four operative interferometric detectors around the world, managed by three scientific collaborations: The LIGO Scientific Collaboration, which runs two interferometers located in the United States of America; the GEO Collaboration, which runs a 600 meter interferometer located in Hannover, Germany, and the Virgo Scientific Collaboration, which runs an interferometer located in Italy. LIGO and Virgo possess the most sensitive instruments, as they are an enhanced version of the initial facilities; regarding GEO, it is mainly used to prototype and test new features which may become useful for future generations of gravitational wave observatories.

Such an amount of observatories allows for a substantial amount of measurements to be performed: If we talk about transient signals, a triple interferometer detection could yield a precise definition of the sky position, as well as accurate measurements of the gravitational wave polarizations; as for continuous signal, it would yield a confident check on its persistence, as well as a confident measurement of its sky location.

1.3.1 Power Spectral Density of a Detector

Interferometric detectors are continuously affected by noise. As we discussed in section 1.2, the expected gravitational wave amplitudes are located at the bleeding edge of current technology; hence, understanding the response of the detector to undesired external perturbations is an important task, both to mitigate them and to prevent spurious results to be treated as detections.

Noise is usually characterized via its *power spectral density* (PSD); that is, the Fourier transform of the autocorrelation of an ensemble of measurements $n(t)$. It is customary to use the *single sided* power spectral density, which is spanned over the positive frequency region of the spectrum, since any performed




Figure 1.5: Amplitude Spectral Density of the LIGO Livingston (left) and Hanford (right) interferometers during the last observation runs, O1 and O2. Source: [22].

measurement is going to be a real quantity:

$$S_n(f) = 2 \int_{-\infty}^{\infty} \langle n(t)n(0) \rangle e^{-2\pi i f t} dt. \quad (1.11)$$

$\langle \cdot \rangle$ denotes ensemble average. The actual estimation of power spectral density depends on the kind of analysis performed, as one has to take into account the presence of real signals within the data: Some pipelines rely on robust estimations based on the statistical median, while others plead for auto-regressive estimations. Equivalently, one could quote the amplitude spectral density (ASD), defined as the square root of the PSD.

Different interferometers will differ on its PSD, either by construction or due to environmental disturbances. As so, a new magnitude is introduced to quote the actual sensitivity to a gravitational wave in a noise-independent fashion. For a gravitational wave with amplitude h_0 , we define its *sensitivity depth* as the ratio between the ASD and the amplitude of the gravitational wave

$$\mathcal{D}(f) = \frac{\sqrt{S_n(f)}}{h_0}. \quad (1.12)$$

The weaker the signal is with respect to the noise, the *deeper* it is burdened into the noise.

Figure 1.5 shows the noise amplitude spectral density of the LIGO detectors during the last observation runs. Multiple effects can cause noise at different frequency ranges. For instance, seismic or instrumental vibrations couple to the low frequency band of the detector, up to 20Hz, while thermal fluctuations on the mirrors or quantum fluctuations coming from the laser beam dominate at higher frequencies. In addition, there is a vast population of transient noise sources related to the environment of the detector, such as adverse meteorological conditions, interactions with the local fauna or human activity.

Searches for transient signals must worry about that kind of noise, since its appearance in the detector resembles the shape of a gravitational wave, due to the general trend followed by transient events, ringing and decaying as a bell would do. However, if we were targeting a *persistent* signal, our concerns would be focused on the ensemble of spikes or *lines* shown in figure 1.5, which correspond to persistent sources of noise. Some of them are widely know, such as the harmonics of 60Hz, which correspond to the frequency of A.C. in the US; others, such as the *comb* of lines presented in both detectors, are poorly understood and require a further analysis. An study on this subject for the last observation runs can be found in [22].




Figure 1.6: Noise contributions for the Livingston observatory during the O2 observing run. Source: [79].

1.3.2 General Noise Contributions

We proceed to expose the main sources of noise in an interferometric gravitational wave detector. Figure 1.6 shows the main contributions at each frequency for the LIGO Livingston interferometer during the O2 observing run.

Below 40Hz, the main noise contribution is given by the *seismic noise*; that is, *vibrations*. Any nearby vibration can potentially reach the detector and couple to the suspension wires of the mirrors. This is not limited to telluric motion, as there are contributions from the surrounding environment, such as gravitational gradients, or direct effects on the interferometers, such as human activity,

The central region of the spectra is dominated by thermal noise: Mirrors have finite temperature, which implies certain vibration of its components. The use of different materials and components can low this contribution. From a more extreme point of view, it is feasible to work under cryogenic temperatures, as its done by the KAGRA interferometer, located in Japan [39].

Finally, above 200Hz the dominant noise comes from the laser, called *shot noise*. Measurements are performed on a photodiode, using the interference of laser. Due to the quantum effects that get into play, there exists an intrinsic uncertainty into such a measurement. Nowadays, the use of *squeezed vacuum* is being proposed as a solution to this particular contribution.

As interferometers progress, so do their sensitivities. Figure 1.7 represents the sensitivity curves of four interferometric detectors. Comparing to the sensitivity of GEO600, the improvement on the Advanced LIGO and Advanced Virgo detectors is clear.




Figure 1.7: Amplitude spectral density of different gravitational wave interferometers using ~ 1 day of data during the O3 observing run. GPS time is shown in the top label. Vertical axis represents ASD and horizontal axis represents frequency, both in logarithmic scale. Each color represents a different interferometric detector. Source: [79].

1.3.3 Persistent Spectral Disturbances: Lines

An important contribution on the noise spectral for the search of continuous gravitational waves are *lines*. They represent persistent spectral disturbances, which appear as highly concentrated amount of power into the power spectral density of a detector. Figures 1.5, 1.6 and 1.7 show a sample of them.

There are several types of lines: Calibration lines are located at certain frequencies on purpose, in order to gauge the behavior of the detector. Power lines are a direct consequence of the AC electricity in the United States; they are located at the harmonics of 60Hz, and can be clearly seen in figure 1.7. *Violin modes* are due to the mechanical suspensions of the mirrors. *Combs* are sets of lines the nature of which is not clearly understood; they were prominent in O1 data, but they seem to be faded away during the O2 observing run, as clearly seen in figure 1.5.

Summarizing, interferometric detectors are affected by a wide variety of noise artifacts, even though technological improvements are being continuously made in order to suppress its contribution to the search for gravitational waves.

Neutron Stars as Gravitational Wave Sources

The following chapter will be devoted to the description of Neutron Stars as continuous gravitational wave sources. The first section will expose such objects from the physical point of view, giving a brief insight into the information they could provide. The second one will describe the continuous signal yielded by neutron stars, as measured from an Earthbound detector. During the third section, we will overview the main data analysis strategies for the detection of such signals. Finally, the fourth section will expose the current state of continuous gravitational wave searches within the LIGO and Virgo scientific Collaboration.

2.1 Physics of a Neutron Star

The first appearance of a neutron star (NS) into Physic's literature was due to Bade and Zwicky [86], who proposed them as the product of a certain kind of supernova processes. Due to the lack of plausible candidates, they befallen into the oblivion until a few decades later, when a team guided by Anthony Hewish detected a regular, stable radio signal coming from a certain sky position [34]. After an initial speculation on extraterrestrial life, the consistency of such a signal with a spinning neutron star pointing one of its material jets towards the Earth in a periodic motion was proven [32]. Such sources, which produced pulsations in the radio frequency band, were called *pulsars*.

Neutron stars are one of the most compact objects in the Universe, condensing around 1 – 2 solar masses (M_{\odot}) into a radius of 20 – 30km. Its main channel of formation is due to the collapse of certain stars with masses above $8M_{\odot}$ [20].

Its internal structure is a nowadays matter of discussion, since it represents one of the most extreme conditions in which matter can be found in the Universe. Current models agree in the existence of an external crust made of metals, which encloses an internal, highly degenerated matter region. The density of the internal fluid is such to produce a high number of electronic capture processes, creating neutrons that can not decay via β -processes due to the degeneracy of the surrounding electrons. This causes a neutron *drip*, which, close to the internal region of the star, becomes a superfluid of neutrons. In addition, nuclei are effectively dismantled, not only allowing protons and electrons to wander around, but also producing an effective proton-proton pairing that yields a *superconducting* fluid. As for the innermost region, some models predict a solid nucleus of quarkonic matter [48].

The pulsating mechanisms by which the first pulsars were observed is also under current discussion. Its existence can be justified in terms of the frozen magnetic field within the superconducting fluid of the

neutron star: Due to the rapid rotation of the neutron star, an electric field is induced in its vicinity, which ends up forming a magnetosphere of particles being dragged and ejected as its velocity is increased.

Glitches are another observed process into pulsars. They consist in an abrupt decrease of the rotation period of the body. Its underlying mechanism is not currently well understood, but it could be related to instabilities into the internal structure.

Being such an extreme case of degenerated matter, any information probes into neutron stars could lead to major improvements on the current understanding of quantum chromodynamics (QCD). However, electromagnetic technologies are only able to detect a certain type of neutron stars: those ones whose jet points towards the Earth; otherwise, no light comes to us. For this precise reason, we aim to describe the possible gravitational wave emission mechanisms of neutron stars, in order to detect its physical paw into the spacetime curvature.

2.1.1 Emission Mechanisms

We proceed to describe the gravitational wave emission mechanisms of *isolated*, rapidly spinning neutron stars. Such celestial bodies can be found in multiple astrophysical situations, such as is binary systems with a companion star, with another neutron star or even with a black hole. Each of those systems presents its own emission mechanisms, which will not be discussed in detail in this work. We will focus on *continuous* emission mechanisms; that is, gravitational wave emissions which persist during long periods of time.

There are three plausible emission mechanisms within the LIGO frequency band. Each of them relies on the sustain of a deviation from axisymmetry into the structure of the spinning neutron star, producing an overall time-dependent quadrupolar moment.

- Crustal deviations from axisymmetry

The first mechanism was already exemplified during section 1.2. It consist in the sustain of a *bump* into the outermost layer (the *crust*) of the neutron star. As mentioned before, such a deformation is quantified in terms of the equatorial ellipticity ϵ , given by (1.6), leading to an characteristic amplitude given by (1.9). The frequency of this gravitational waves would be twice the spinning frequency of the neutron star; that is, $f_0 = 2\nu$ using the previously established notation.

In order to fully understand the magnitude of such deformations, one recurs to numerical descriptions of the crust, using a wide variety of hypothesis regarding its structure and composition. For a standard Neutron Star, maximal ellipticity values are given by $\epsilon \sim \mathcal{O}(10^{-6})$, although lower values are equally likely due to the high uncertainty on the actual crust composition. Ellipticity can grow until $\epsilon \sim \mathcal{O}(10^{-3})$ if we include quark matter into the equation of state [37, 42].

- Non-axisymmetric fluid instabilities

A second mechanism relies on internal structure of neutrons stars: Coriolis force could lead to the excitation of toroidal *r-modes* into the internal fluid, leading to sustained, non-axisymmetric instabilities. This particular mechanism relies on young neutron stars, or even neutron stars in an accreting ambient, to satisfy the suitable conditions into the surrounding fluid.

Nevertheless, the nature of *r-modes* tends to be futile, lasting for too short periods of time as to be considered into continuous gravitational wave searches [54].

- Free precession

The dynamics of a neutron star present a phenomenon of precession when the angular momentum of the star is not aligned with its symmetry axis. This situation could be given by several mechanisms: a glitch




Figure 2.1: Population of known pulsars in terms of its rotation period and its first derivative. Pulsars in a binary system are highlighted by open circles. Lines of constant magnetic field (dashed), characteristic age (dot-dashed) and spin-down energy loss are also shown. The vertical green line marks the neutron stars within the LIGO frequency band. Source: [51].

process undergoing into the celestial body, a nearby encounter with another neutron star, or even the presence of a crustal imperfection.

The amplitude of the resulting gravitational wave is described in terms of the wobble angle of the precession θ_w as

$$h_0 \sim 10^{-27} \left(\frac{\theta_w}{0.1} \right) \left(\frac{1 \text{ kpc}}{d} \right) \left(\frac{\nu}{500 \text{ Hz}} \right)^2, \quad (2.1)$$

where d represents the distance to the source and ν the rotation frequency of the star.

The possibility of detecting this type of signal was discussed by [44], leading to rather pessimistic results due to the damping mechanisms which affect the wobble angle.

2.1.2 Population of Neutron Stars

Estimations on the neutron star population of the galaxy yield close to 10^8 neutron stars in the Milky Way, taking into account the rate of stars going supernova in our galaxy [51]. The actual estimation breaks this number into 10^5 pulsars (with only ~ 2700 discovered) and 10^7 unseen neutron stars, either because its faintness or the absence of electromagnetic jet pointing towards the Earth.

Figure 2.1 shows the current known distribution of neutron stars in terms of its rotation period and its first derivative¹. The decrease in the rotation period of a neutron star is a natural consequence of its

¹[35] maintains a census of known neutron stars, as well as its current values of rotating frequency and spin-down.

behavior, since electromagnetic waves are continuously expelled due the strongly varying electromagnetic fields, as well as possibly gravitational waves, as stated before. For later convenience, we will equivalently use the rotation frequency of the star ν and its first temporal derivative $\dot{\nu}$, called *spin-down*, in terms of the rotation period.

The period distribution suggest the existence of, at least, two great groups of pulsars; namely, the ones with periods between $0.25s$ and $2s$ and the ones with periods around milliseconds.

The most well-known pulsars, Vela (PSR B0833-45) and the Crab (PSR B0531+21), belong to the first group. They established a solid, positive evidence in favor to the relation between pulsars and supernovas.

As for the second group, the fact that they predominantly belong to binary systems, together with its high rotation frequencies, suggests that they are in fact *recycled* neutron stars, spun up by the accretion of matter [55]. The accretion process could produced sustained quadrupole deformations which may yield an emission of gravitational waves [83]. It is observed that this sub-family of pulsars rotates at a low frequency, compared to the maximum imposed by centrifugal break-down. A proposed mechanism which allows this kind of behavior is the emission of continuous gravitational waves, which compensates the accretion spin-up.

A further pulsar characterization can be made in terms of the *braking index* n , which characterizes the main source of energy loss of the neutron star by means of the relation between its frequency and its spin-down

$$\dot{\nu} = -k\nu^n, \quad (2.2)$$

where k is a structural constant of the star. Differentiating with respect to time, we obtain a measurement for n

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2}. \quad (2.3)$$

A braking index of $n = 3$ corresponds to a pure magnetic brake by means of dipolar radiation; pure gravitational brake due to sustained quadrupole deformations are characterized by $n = 5$, while saturated r-mode emission implies $n = 7$. Values below $n = 3$ can be justified in terms of time dependent magnetic dipolar orientations or outflowing particles.

This kind of information is useful to constraint the amount of gravitational waves radiated by neutron stars. Currently, only nine pulsars have measured braking indexes due to multiple effects disturbing a clear measurement of $\ddot{\nu}$, ranging from $n = 0.9 \pm 0.2$ to $n = 3.15 \pm 0.03$ [52].

A second characterization can be made in terms of the *maximum* amount of energy released as gravitational waves by a neutron star, imposing the current spin-down of the star to depend entirely on its gravitational wave emission. Let us work out the case for a neutron star with a crustal deformation characterized by a certain equatorial ellipticity ϵ . As stated during section 1.2, the gravitational wave luminosity of such a physical system is given as

$$\mathcal{L}_{GW} = \frac{1}{10} \frac{G}{c^5} (\pi\nu)^2 I_3^2 \epsilon^2, \quad (2.4)$$

where ν is measured in Hertz, hence the π factors. As for the rotational kinetic energy of the neutron star it is given by [57]

$$K_{rot} = \frac{1}{2} I_3 (2\pi\nu)^2. \quad (2.5)$$

At this point, it suffices to set $\mathcal{L}_{GW} = -\dot{K}_{rot}$ in order to set upper bounds on the desired magnitude. If we assume $\dot{I}_3 = 0$, which corresponds to a stable neutron star, we obtain an upper bound on the equatorial

ellipticity given by

$$\epsilon^{sd} = \sqrt{\frac{5c^5}{512\pi^4 G I_3} \frac{|\dot{\nu}|}{\nu^5}}. \quad (2.6)$$

Following (1.9), the corresponding gravitational wave amplitude is given by

$$h_0^{sd} = \frac{1}{r} \sqrt{\frac{5GI_3}{2c^3} \frac{|\dot{\nu}|}{\nu^5}}. \quad (2.7)$$

Section 2.4 will show the utility of this magnitude during: Even though no gravitational wave signals have been detected so far, the analysis of multiple pulsars can set upper bounds in the actual amplitude of gravitational wave emissions. Comparing this upper bounds with the corresponding spin-down limit of each pulsar yields valuable information about the physics governing that particular neutron star, as well as about the current state of gravitational wave detectors and search algorithms.

2.2 Signal from an Isolated Neutron Star

Independently of the multiple mechanisms which could sustain a continuous gravitational wave emission from a neutron star, the main features of the signal will be completely identical: A *quasi-monochromatic* emission, slowly damped due to the multiple radiation channels of the astrophysical body. It is worth noting that the neutron star could be surrounded by an accreting disk of matter which, effectively, spins it *up*, producing a net increase in its rotation frequency.

The description of a continuous gravitational wave signal depends on the referential frame being used, either the proper frame of the source, the solar system baricenter (SSB) or the detector's frame. We will start from the frame of the source, explicitly describing the propagation of the signal until its arrival to the interferometric detector as stated by [41].

2.2.1 Phase and Frequency: From the Source to the Detector

The phase of a gravitational wave signal from the referential frame of the source is described, in terms of its proper time τ , as

$$\Phi_{\text{NS}}(\tau) = \phi_0 + 2\pi \sum_{n=0}^s \frac{f_n^{\text{NS}}}{(n+1)!} \tau^{n+1}; \quad (2.8)$$

consequently, the instantaneous frequency of the gravitational wave signal is given by its time derivative

$$f_{\text{NS}}(\tau) = \frac{1}{2\pi} \frac{d\Phi_{\text{NS}}(\tau)}{d\tau} = \sum_{n=0}^s \frac{f_n^{\text{NS}}}{n!} \tau^n, \quad (2.9)$$

which clarifies the meaning of the expansion coefficients: f_0^{NS} represents the initial frequency of the signal at certain $\tau = 0$, whereas $f_{n>0}^{\text{NS}}$ model the spin-down of the signal using a set of s polynomial terms.

The propagation of this signal into the detector's frame will be done in two steps: The first one will put the signal into the SSB, and the second one will push it forward to the detector. Doing so allows for a separation between two families of effects, namely The ones related to the propagation of the signal from the source, and the ones related to the proper movements of the Earth around the SSB and its own axis.

We start by assuming an isolated neutron star moving in a uniform way with respect to the SSB, although one could drop this assumption at the price of taking further parameters into account. Defining t_{SSB} as

the time coordinate of the solar system baricenter, the condition which links the phase measured by both referential systems is given by

$$\Phi_{\text{SSB}}(t_{\text{SSB}}) = \Phi_{\text{NS}}(\tau(t_{\text{SSB}})), \quad (2.10)$$

which takes into account the propagation time of the signal as well as the relativistic corrections to change between time coordinates in relative motion. Applying the previous assumption on the proper movement of the star, one ends up with a power expansion analogous to (2.8) for the phase into the SSB

$$\Phi_{\text{SSB}}(t_{\text{SSB}}) = \Phi_0 + 2\pi \sum_{n=0}^s \frac{f_n}{(n+1)!} t_{\text{SSB}}^{(n+1)}, \quad (2.11)$$

where the exact form of $t_{\text{SSB}} = t_{\text{SSB}}(\tau)$ has been kept implicit. It is important to note that the spin-down coefficients described in the solar system baricenter do not have to be equal to the ones in the proper frame of the source, $f_n^{\text{NS}} \neq f_n$. Consequently, the frequency observed from the SSB will be given by

$$f_{\text{SSB}}(t_{\text{SSB}}) = \frac{1}{2\pi} \frac{d\Phi_{\text{SSB}}}{dt_{\text{SSB}}} = \sum_{n=0}^s \frac{f_n}{n!} t_{\text{SSB}}^n. \quad (2.12)$$

Last, we need to push the signal forward to the detector frame. In order to do so, let us define \vec{v} as the velocity vector of the detector in the SSB frame, \vec{n} as the sky position of the source *with respect to the detector*, and $\Delta\vec{r} = \vec{r} - \vec{r}_0$ as the displacement of the detector in the SSB frame with respect to a fiducial initial position \vec{r}_0 .

We define t as the time coordinate of the detector frame. This new time can be related to t_{SSB} by taking into account several effects in its relativistic formulation, such as the Shapiro effect [66], which is due to the deflection of null geodesics in the presence of matter, and the Doppler effect, which is produced by relative motion among wave emitters and receptors.

For the former, it can be shown that its effects become unobservable for the matters of continuous gravitational waves [41]; as for the later, the characteristic speed of rotation and translation of the Earth are at most $10^{-4}c$ and $10^{-6}c$ respectively, hence relativistic corrections can be dropped from the expansion.

Therefore, the frequency of the gravitational wave signal at the detector frame is expressed as

$$f(t) = f_{\text{SSB}}(t) \left(1 + \frac{\vec{v}(t) \cdot \vec{n}(t)}{c} \right), \quad (2.13)$$

where the detector's time t is related to t_{SSB} accordingly with the stated simplifications

$$t_{\text{SSB}} = t - t_0 + \frac{\Delta r(t) \cdot \vec{n}(t)}{c}, \quad (2.14)$$

being t_0 a fiducial initial time.

2.2.2 Amplitude of the Gravitational Wave

The response of an interferometric detector to a passing gravitational wave was already discussed in section 1.3: The measured strain depends on the antenna response functions of the detector and the gravitational wave polarizations of the signal, which depend on the actual source being studied.

For the case of a continuous wave signal, they can be stated in terms of the signal phase $\Phi(\tau)$ and two amplitude functions $A_{+,\times}$ as

$$h_+(\tau) = A_+(\tau) \cos \Phi(\tau), \quad (2.15)$$

$$h_{\times}(\tau) = A_{\times}(\tau) \sin \Phi(\tau), \quad (2.16)$$

where the notation is consistent with the one of the previous section. For a source like the deformed neutron star described in (1.9), such functions are

$$A_{+} = \frac{1}{2} h_0 (1 + \cos^2 \iota), \quad (2.17)$$

$$A_{\times} = h_0 \cos \iota, \quad (2.18)$$

where ι represents the angle between the neutron star's axis of rotation and the direction of \vec{n} . $\iota = 0, \pi$ corresponds to a circularly polarized wave, while $\iota = \pi/2$ corresponds to a linearly polarized wave. As clearly seen from the previous expression, they represent the best and worst cases of a signal, respectively.

2.3 Data Analysis for a Continuous Gravitational Waves

As discussed during section 1.2, continuous gravitational waves represent a technological challenge due to its faintness, in comparison with other forms of gravitational waves emission. Even taking into account the increasing sensitivity of the current LIGO and Virgo detectors, one needs to come up with the suitable data analysis strategy in order to recover a signal buried into the noise. To do so, several frameworks have been developed, depending on the characteristics of the source being searched, and the way information is treated during the process:

- Targeted Searches

As suggested by its name, this type of search deals with sources whose parameters are already restricted to a tight region of the parameter space thanks to a complementary information channel. For instance, we could dispose the sky position of a pulsar via electromagnetic observations, or even describe its actual phase evolution taking into account glitches and any observed transient effect. This situation eases the use of highly significant methods under controlled computational budgets with respect to other kinds of searches.

- Blind Searches

On the contrary, this searches look for a priori unknown continuous wave sources. The strategy is quite different with respect to the previous case, since the lack of constraints forces us to perform wide analysis on the parameter space. As a consequence, it is not affordable to attempt the most sensible strategies due to its high computational cost. Hence, the sensitivity is lower with respect to targeted searches.

- Directed Searches

In some cases, there are certain sky regions in which neutron stars are more likely to be found, even when no electromagnetic counterpart has been detected. Such is the case for supernovae remnants, which can be well located in the sky. This searches focus on the frequency and spin-down parameters of the source, conforming an intermediate step between targeted and blind searches.

This division already introduces the need for different data analysis strategies: On the one hand, we require powerful methods, able to analyze specific regions of the parameter space with a high sensitivity; on the other, it will be useful to dispose of less powerful strategies in order to swap wide domains of the parameter space and spot interesting regions.

We will describe the two main lines of analysis of gravitational wave signals, *coherent* and *semicoherent* methods. Its difference relies on the way they treat the data, with direct implications on its computational costs.

2.3.1 Coherent Methods

The basic data analysis procedure consist of a direct comparison between data and signals; such an algorithm is called *matched filtering*, and its derivation can be statistically formulated in terms of a maximum likelihood estimator. The performance of this algorithm relies on the capability to produce *templates* of continuous wave signals for different combinations of parameters, being the dependency of such templates with respect to its parameters the key feature on the computational costs of a search.

From an operational point of view, it is said that this algorithm takes into account the *phase* of a signal, which provides more information than its frequency. Hence it is labeled as a *coherent* method.

Let us consider a parameter space of templates described by four parameters $\vec{\lambda} = (f_0, \vec{n}, f_1)$, where f_0 stands for the frequency of a signal, f_1 for one the spin-down parameter and \vec{n} for the sky position. Given that we need to compare multiple templates in order to unveil the presence of a signal into the data, it is natural to ask *how many* templates should one take into account for this purpose.

As described by [61], the answer depends on the frequency of the template, due to the Doppler modulation suffered by the signal. If we consider a small volume of the parameter space, say $d\vec{\lambda}$, the amount of templates within that volume will be given by

$$dN_t^c \propto T_{\text{obs}}^5 f^2 d\vec{\lambda}, \quad (2.19)$$

where T_{obs} stands for the temporal duration of the data. Moreover, each templated has to be integrated by the matched filtering procedure during the whole span of the data; hence, the computational cost for an amount of dN_t templates is going to be

$$dC^c = T_{\text{obs}} dN_t^c \propto T_{\text{obs}}^6 f^2 d\vec{\lambda}. \quad (2.20)$$

As described in [46], this analysis easily overcomes the computational costs affordable by current computers for wide searches of the parameter space. However, it is well suited for targeted searches, due to the reduced amount of templates to take into account.

2.3.2 Semicoherent methods

Semicoherent methods constitute the second major family of data analysis algorithms. Opposed to the previous one, the main source of information is not the phase of a signal, but its frequency. The main strategy consists in the split the whole data, with a duration of T_{obs} , into smaller segments, with a duration $T_{\text{coh}} < T_{\text{obs}}$. Each segment is analyzed using the previously mentioned *coherent* method, and the results are added up following a statistical prescription which takes into account the frequency evolution of the signal, hence *incoherently*. As a result, the amount of templates to be taken into account is given by [61]

$$dN_t^s \propto T_{\text{obs}} T_{\text{coh}}^4 f^2 d\vec{\lambda}, \quad (2.21)$$

while the computational costs scales as

$$dC^s = T_{\text{obs}} dN_t^s \propto T_{\text{obs}}^2 T_{\text{coh}}^4 f^2 d\vec{\lambda}. \quad (2.22)$$

Comparing this result with that of a coherent method (2.20), it turns out that

$$\frac{dC^s}{dC^c} \propto \left(\frac{T_{\text{coh}}}{T_{\text{obs}}} \right)^4. \quad (2.23)$$

For a typical segment length $T_{\text{coh}} = 1800s$ and 30 days of observation time, the computational cost reduction achieved by a semicoherent method is $\mathcal{O}(10^{-13})$ with respect to a coherent one.

As a result, the main strategy followed by the LIGO and Virgo collaboration is that of a hierarchy: Interesting candidates yielded by semicoherent methods are further studied under the lens of a coherent method, thanks to the reduction of the interesting parameter space region yielded by the first method.

2.4 State of the Art on Continuous Wave Searches

We proceed to end this chapter with a brief review on the latest results given by the last observing run, O2, performed by the LIGO and Virgo interferometers.

Given that no continuous gravitational waves have been detected so far, the results are quoted in terms of *upper limits* on the gravitational wave amplitude h_0 for a certain confidence level, the exact meaning of which depends entirely on the statistical prescription used by the search. Moreover, upper limits can be translated to another quantities, such as the ellipticity ϵ , by accepting a certain emission mechanism. Finally, searches for known pulsars can also benefit from the previously discussed spin-down limit (2.7), which will ease the task of describing the energy budget of a star.

2.4.1 All Sky Search

The all sky search is devoted to the detection of continuous gravitational waves coming from unknown sources. It is composed by three semicoherent pipelines: *Time-Domain \mathcal{F} -Statistic*, *SkyHough* and *FrequencyHough* [9]; the first one is based in a semi-coherent implementation of the \mathcal{F} -statistic, which is the optimal test for the detection of continuous gravitational waves using a coherent analysis; the later ones are different implementations of the Hough Transform, to be discussed during the following chapter.

Its results are given in terms of frequentist upper limits on the amplitude h_0 . Its construction is based in the generation of a random ensemble of continuous wave signals. Then, the ensemble is added to the detector noise, in order to be looked for such. The amount of detected signals yields an estimation on the minimal strength for a signal to be detected under the current configuration of the searches.

The upper plot in figure 2.2 shows the 95% confidence upper limits for the three pipelines. In addition, the lower left and lower right panels propagate the amplitude upper limit to the ellipticity and the first spin-down parameter for different distances. The diverse morphology of each curve is due to the details in the upper limit computation of each pipeline.

These values are useful in the sense that they provide astrophysical information on the nature of neutron stars. For instance, the upper limits on ellipticity are *below* the maximum ellipticity expected for a neutron star using a standard equation of state [42], whilst the upper limits on spin-down are to be compared with the expected energy emission of the sources.

2.4.2 Targeted Search

The targeted search looks for gravitational wave emission from 222 known pulsar using three different search methods, namely a *Bayesian* parameter inference using time-domain data, a coherent \mathcal{F} -statistic and the *5n-vector* method [76].

The pipelines benefit from optical observations of the pulsars to restrict the parameter space and legitimate a coherent approach to the data analysis. The construction of upper limits is different for each pipeline, since their statistical development is based in different paradigmata. The results are shown in figure 2.3

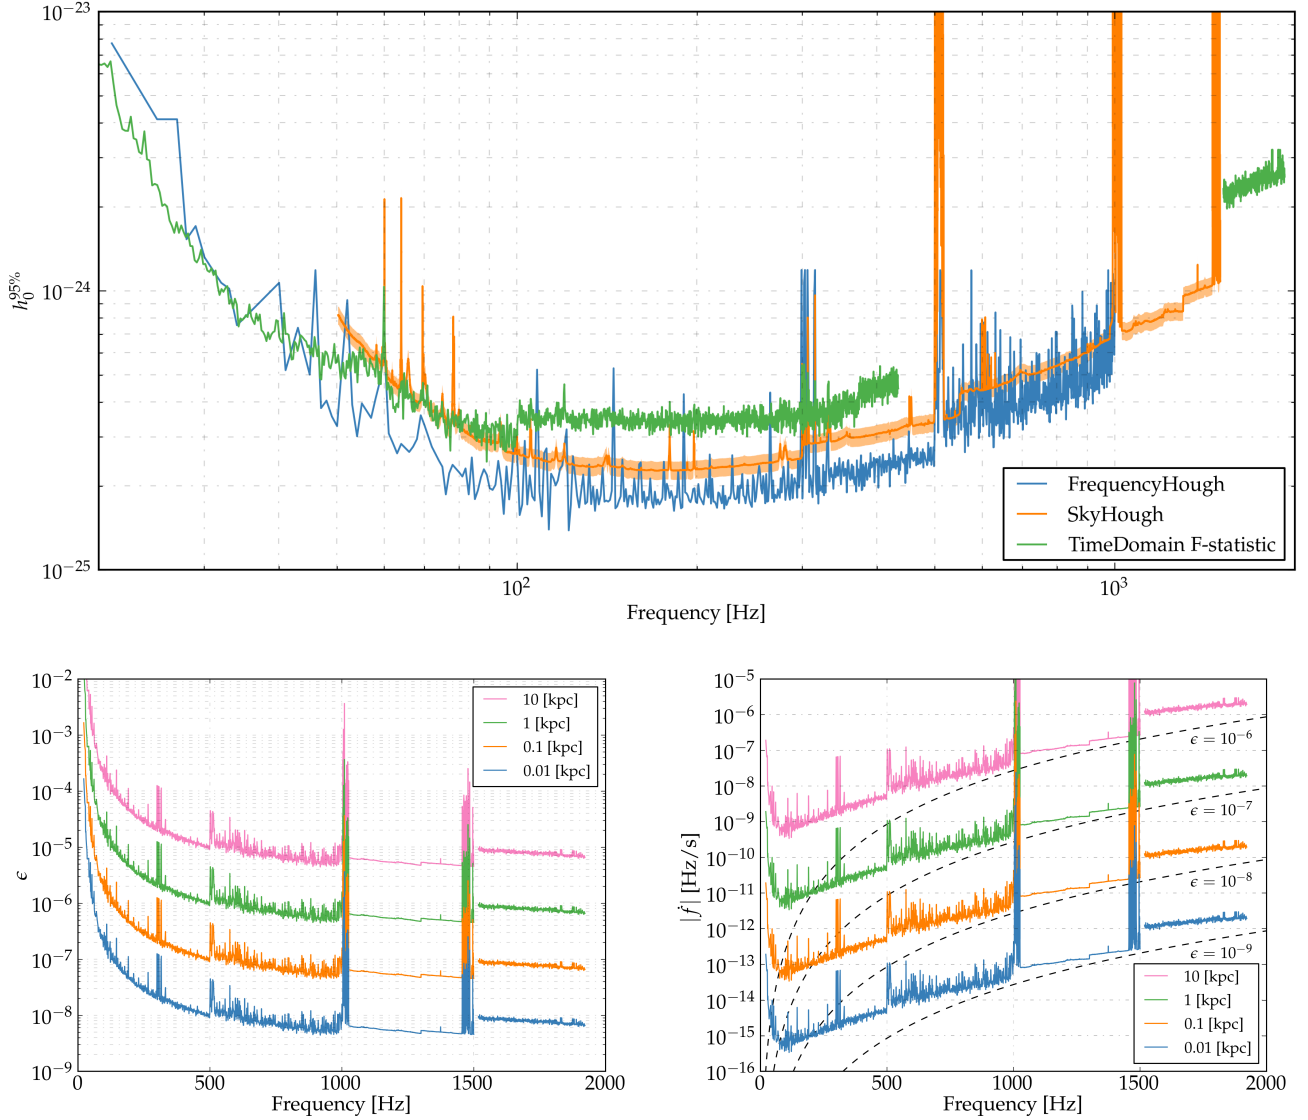


Figure 2.2: Results of the All Sky search performed over O2 data. **Upper:** 95% confidence frequentist upper limits on the gravitational wave amplitude for the three pipelines involved in the search. **Lower Left:** Upper limits propagated to equatorial ellipticity using equation (1.9) for the canonical value $I_3 = 10^{38} \text{ kg} \cdot \text{m}^2$ and different distances. **Lower Right:** Upper limits propagated to the maximum allowed spin-down value for a spinning neutron star assuming a pure gravitational wave emission. Dashed lines represent regions of constant source ellipticity. Source: [9].

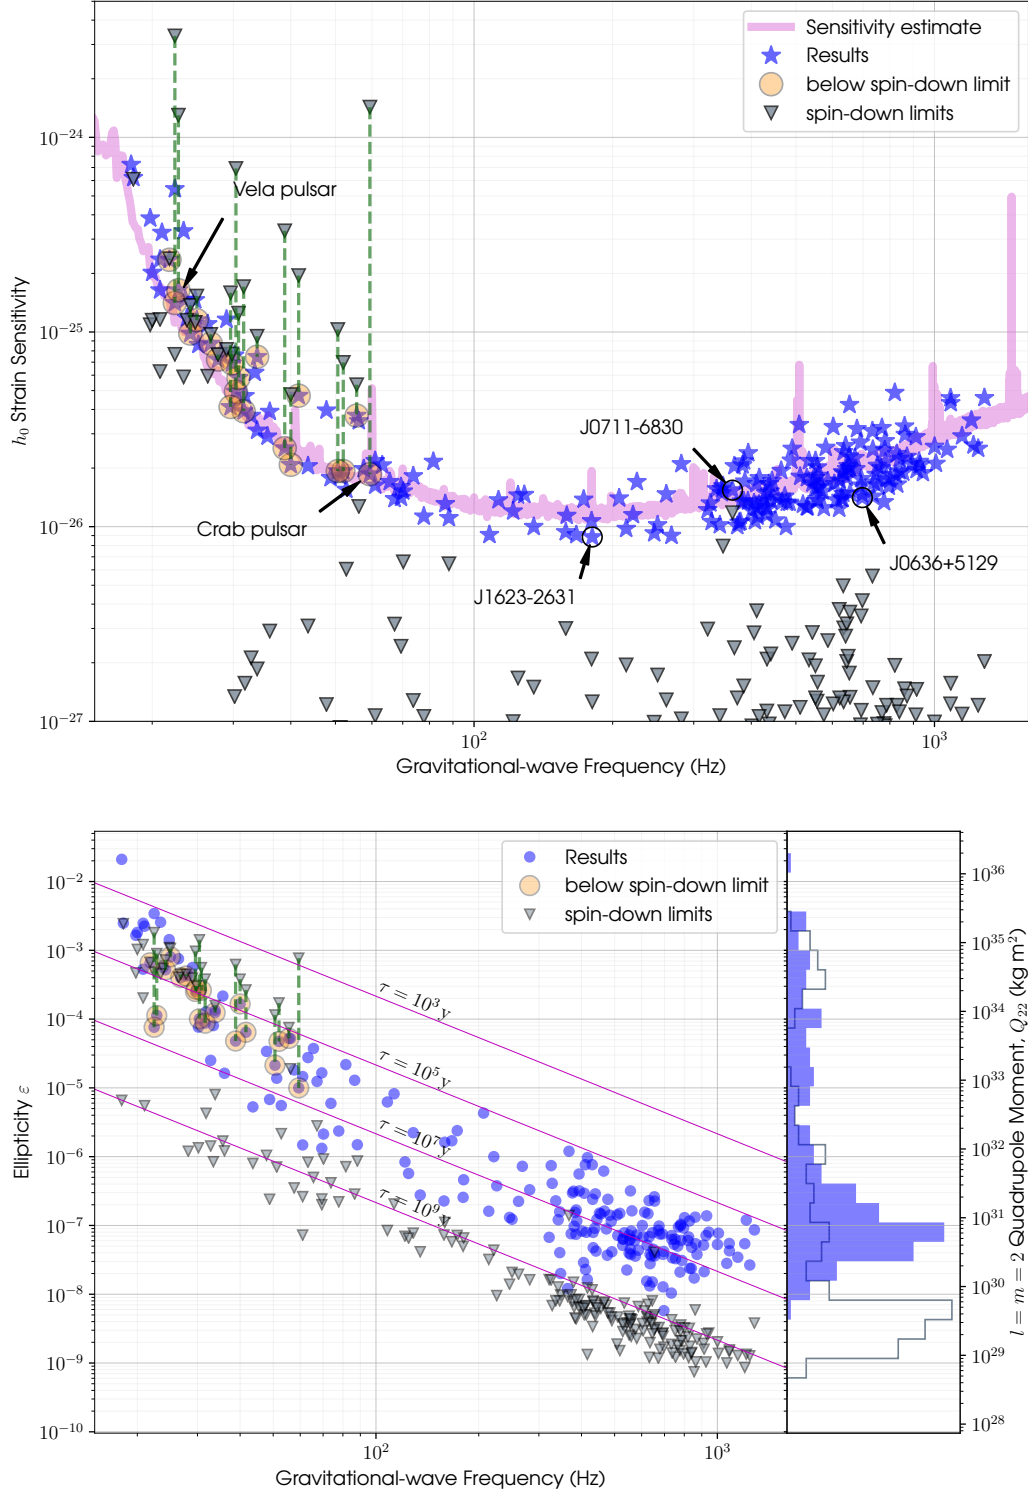


Figure 2.3: Results of the Targeted Search after the O2 observing run. **Upper:** Amplitude upper limits for the known pulsars, compared to its spin-down limit. **Lower:** Propagation of the upper limits to the maximum allowed ellipticity for the known pulsars, as well as a measurement on its quadrupole mass moment. The conversion to ellipticity is alike to the one made in figure 2.2. Source: [76].




Figure 2.4: 95% frequentist upper limits on the gravitational wave amplitude coming from a set of 33 pulsars with known sky position. Pulsars that glitched during the observing run got assigned two upper limits: one *before* the glitch (BG), and one *after* the glitch (AG). For comparison, spin-down limits are quoted with 1σ error bar on the distance uncertainty. Source: [75].

Dealing with known sources eases the comparison of results with the spin-down limits derived during previous sections, imposing tight restrictions on the energy budget of each neutron star.

In this case, the spin-down limit is beaten for several pulsars, being Vela and the Crab, long-run significant targets of this type of searches, among them. Respectively, the energy budget obtained by their upper limits places the gravitational wave emission *below* the 0.017% and 0.18% of the overall energy emission.

Under the light of this results, is expected for the O3 observing run to yield measurements bellow the spin-down limit for millisecond pulsars, which represent another prominent target on the search for continuous gravitational waves.

2.4.3 Narrow Band Search

The narrow band search enters the domain of targeted searches, as it is designed to look for gravitational wave signals coming from neutron stars whose position is well known via a third-party procedure, usually an electromagnetic measurement. However, it uses a coherent method which allows for a certain mismatch between the signal model and the phase measured by electromagnetic observations, the *5-vector NarrowBand* pipeline [75]. This phenomenon can be produced by multiple factors, such as electromagnetic torques, free precession, or inaccurate sky position measurements due to gravitational interaction

Figure 2.4 shows the 95% frequentists upper limits on amplitude given by the search after the O2 analysis. The mechanism to generate upper limits follows the philosophy of all sky searches, although it has to be accommodated to the particular gimmicks of the narrow band analysis. Those neutron stars which




Figure 2.5: Results of the Directed Search for one of the sky position of *Fomalhaut b* after the O1 observing run. The horizontal line represents a first guess on the upper limit using energy conservation arguments. Each point represents an upper limit on a 1Hz frequency band. *Wide* and *deep* refer to the best and worst case scenarios for the energy conservation argument. Points located in an upper line represent 1Hz bands where data quality allowed no upper limit to be set. Source: [10]

undergone through a glitch process during the observing run were analyzed twice: before and after the glitch.

The minimum bound on gravitational wave energy loss imposed by the upper limits is 0.8%, while the maximum ellipticity is $\epsilon \sim \mathcal{O}(10^{-5})$. It is worth mentioning that the upper limits over pulsars undergoing a glitch process are worse than the ones obtained by the O1 analysis due to the separated analysis between *before-glitch* data and *after-glitch* data.

2.4.4 Directed Search

Directed searches select a certain number of interesting sky regions where a neutron star could be located. During the last directed search [10], sixteen candidates were selected: fifteen supernova remnants and *Fomalhaut b*, a directly imaged extrasolar candidate which could turn out to be an old neutron star [36].

The search is performed using the multi-interferometer \mathcal{F} -statistic, which is the optimal statistic for the search of continuous gravitational waves [41]. It looks for three parameters, namely frequency and two spin-down parameters. As already commented, sky locations are obtained by other means, such as electromagnetic observations.

Upper limits for the *Fomalhaut b* search are shown in figure 2.5. They are set in a similar manner to the one used for all sky searches. Equivalently, the maximum allowed ellipticity is $\epsilon \sim \mathcal{O}(10^{-9})$ for the high frequency range. The object is too old to sustain an r-mode emission. The same celestial body was looked for by the *Chandra* X-ray telescope, although its emission was not detected [68]. All in all, the true nature of *Fomalhaut b* is unclear, yet the possibility of being a nearby old neutron star is not ruled out by any of the available channels of information.

2.5 Summary

Continuous gravitational waves have not been detected so far. Yet, its physical implications contribute to the actual neutron star physics by setting experimental bounds on the theoretical, physical values, which present a high degree of dependency on the available equations of state.

The continuously increasing sensitivity of the detectors allows for the establishment of more stringent upper limits on the multiple searches. In particular, targeted searches towards millisecond pulsars are close to beat its spin-down limit, the first milestone towards the characterization of the energy budget of a pulsar. It is expected that future science runs will deliver the required sensitivity to do so.

As detectors improve, so do data analysis pipelines. Pipelines are in continuous motion, improving multiple steps both in terms of sensitivity and computational efficiency. Different ways of dealing with signal model imperfections or background fluctuations are the main key difference; understanding the their effects on each particular statistic will ease the development of vetoes and follow up strategies, refining the search capabilities towards a first detection. Moreover, the use of Machine Learning algorithms [16, 28], firstly implemented for the recognition of noise artifacts and transient signals, can be extended to the search for continuous gravitational waves. Their pattern recognition capabilities could be useful to neglect spurious artifacts according to its shape into the parameter space. In addition, significant efforts are being put into the implementation of effective computational subroutines like, for instance, the use of Graphical Processing Units (GPU) to optimize particular search loops over the parameter space [23].

Finally, we focused on rapidly-spinning isolated neutron stars as continuous gravitational wave sources; many other celestial structures are able to gravitational waves in a sustained fashion. As shown in figure 2.1, rapidly-spinning neutron stars can be in a binary system, introducing yet another Doppler modulation into the signal. Several search pipelines have been proposed to cover such sources [31], although no gravitational waves have been detected so far.

The Hough Transform for Continuous Gravitational Waves Searches

The Hough transform conforms the computational core of one of the current pipelines on the hunt for continuous gravitational waves within the LIGO Collaboration: The *SkyHough* pipeline. Although its origins can be trace back to a pattern recognition algorithm for particle traces in bubble chambers [38], further developments increased its range of applicability into the image recognition world [27], allowing for the detection of arbitrary shapes.

Within the field of gravitational wave astronomy, the Hough method started its path during the second science run of the initial LIGO detectors (S2) [4], targeting isolated, rapidly-spinning neutron stars, and it kept working through the S4, S5 and S6 science runs [71, 1, 77]. It has been currently used to analyze the observation runs of the advanced detectors, O1 and O2 [6, 8, 9], yielding competitive results with respect to other pipelines.

The modern SkyHough pipeline is built on top of an initial implementation of the Hough transform for the detection of continuous gravitational waves [46], which has been further modified in order to improve its sensitivity. Following its example, several other pipelines were implemented on the same basis, targeting different ideas: The *FrequencyHough* pipeline [13], focusing on continuous waves from isolated sources as well; the *Adaptative Transient Hough* method [59], conceived as an implementation to search for long-duration gravitational wave emission from post-merger remnants; and the *Binary SkyHough* pipeline [23], intended to detect continuous gravitational waves from rapidly-spinning neutron stars in binary systems.

Nowadays, the pipeline has a hierarchical structure [25], which allows for a re-analysis of interesting portions of the parameter space using different tools, such as a χ^2 veto procedure [24]; as a second step, candidates to gravitational wave signals are compared among different detectors; surviving candidates, if they exist, are clustered together, yielding a portion of the parameter space to be followed up using more powerful methods [8, 9].

Such strategies were tested against several other pipelines using a Mock Data Challenge over S6 data [80], which yielded the SkyHough pipeline as the most robust one against noise disturbances.

This chapter will be structured as follows: Section 3.1 introduces the core procedure of the *SkyHough* pipeline; section 3.2 describes the actual implementation of the search, regarding the choice of technical parameters; section 3.3 exposes the statistics of the search, which are fundamental to understand the delivered results; section 3.4 shows the coincidence and clustering post processing steps.

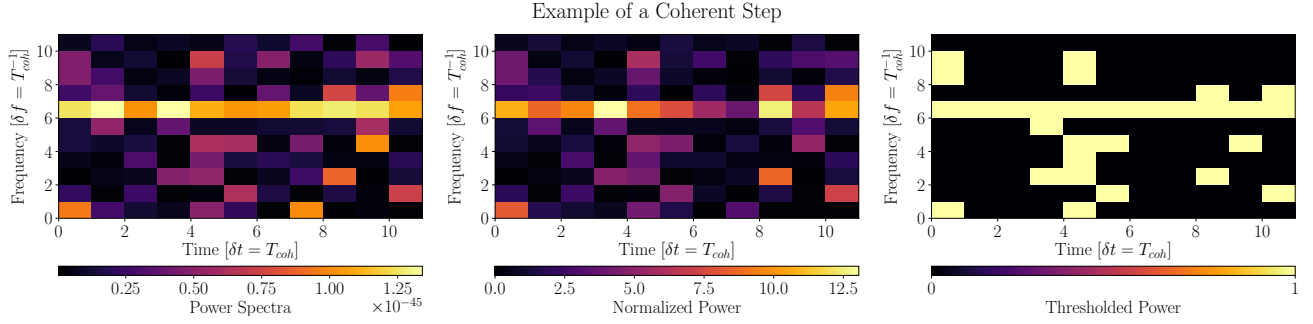


Figure 3.1: Example of a coherent step for a continuous monochromatic signal surrounded by noise. The horizontal axis represents the time of measurement in units of T_{coh} , which is equivalent to time stamp indexes starting from $a = 0$. The vertical axis represents the frequency components of the signal in units of frequency resolution $\delta f = T_{\text{coh}}^{-1}$. **Left:** Power spectra of each coherent segment of data. **Center:** Normalized power of each coherent segment of data. **Right:** Binary map after applying a certain threshold to the normalized power.

3.1 Overview of the Standard SkyHough

The SkyHough pipeline performs a search for continuous gravitational waves using a semi-coherent strategy. We will devote the current section to give an overview of its main formulation, as well as to introduce the internal degrees of freedom of the algorithm in order to focus on the statistical arguments during forthcoming sections.

The method here described attempts to find a continuous gravitational wave signal from an isolated neutron whose frequency evolution matches the pattern produced by the Doppler shift and the spin-down as stated in (2.12) and (2.13); hence, we will focus our attention to the parameters involved in such effects, namely

$$\vec{\lambda} \equiv (f_0, \{f_s\}, \vec{n}) \in \mathbb{R}^+ \times \mathbb{R}^s \times \mathbb{S}^2. \quad (3.1)$$

Each set $\vec{\lambda}$ will be referred to as *template*, being the *parameter space* the set of all possible templates according to (3.1).

The initial step, so called *coherent*, splits the data in multiple time segments with a certain duration T_{coh} . Then, data is processed using the optimal statistic according to the Neyman-Pearson criterion: For a general case, it is given by the \mathcal{F} -statistic [41], which *demodulates* data, removing frequency drifts due to the Doppler shift and spin-down of the source; alternatively, if we restrict T_{coh} to confine the frequency drift of a signal within a certain range, power spectral density becomes optimal statistic in the Neyman-Pearson sense as well. The actual condition over T_{coh} will be discussed during section 3.2.

The second step, labeled as *incoherent*, implements the Hough transform to recognize the trace of a signal within the computed statistic. Thanks to the positional dependency of the Doppler shift, each trace can be related to a certain template, effectively conforming a map from data to the parameter space.

3.1.1 Coherent Step

Consider the *discrete* time-series yielded by a detector $x[j]$, where j labels the j -th measurement. We start by splitting the data into N different segments, starting each one of them at time $t_a, a = 0, \dots, (N - 1)$. The start time of each segment will be referred to as the *time stamp*, and each segment will cover a time span of T_{coh} or, equivalently, M measurements, being $t_j = t_a + j\Delta t$ for each data segment, with $\Delta t = T_{\text{coh}}/M$.

The coherent step will treat each segment separately, computing its Discrete Fourier Transform as

$$\tilde{x}[k] = \Delta t \sum_{j=0}^{M-1} x[j] e^{-2\pi i j k / M}. \quad (3.2)$$

Due to the discrete sampling performed by the measurement process, resolvable frequencies are quantized as $f_k = k\delta f$ with $k = 0, \dots, \lfloor M/2 \rfloor$ and $\delta f = T_{\text{coh}}^{-1}$, while negative frequencies components are easily related to the positive ones thanks to the fact that the Fourier transform of a real function is Hermitian.

For each segment, we compute its power density $|\tilde{x}[k]|^2$ and normalize it to the background noise power, defining the *normalized power*

$$\rho[k] = \frac{|\tilde{x}[k]|^2}{\langle |\tilde{n}[k]|^2 \rangle}, \quad (3.3)$$

where $\langle \cdot \rangle$ represent an ensemble average over multiple realizations of the noise. As previously stated, the noise power spectra $\langle |\tilde{n}_k|^2 \rangle$ can be related to the single-sided power spectral density $S_n(f)$, yielding

$$\rho[k] \simeq \frac{2 |\tilde{x}[k]|^2}{T_{\text{coh}} S_n(f_k)}. \quad (3.4)$$

The estimation of $S_n(f)$ depends on the implementation of the pipeline. *SkyHough* uses a running-median estimation, which prevents major effects due outliers; on the other hand, the *FrequencyHough* implementation [13] uses an autoregressive algorithm.

Then, we digitize the power spectra into a binary map consisting of ones and zeros using power threshold ρ_{th} : Everything above the threshold will get assigned a one, while everything below will get assigned a zero. The definition of such a threshold will be explicitly stated in a further section, using statistical arguments.

The result of this step is a discretized time-frequency plane, where the time domain is discretized according to the data segments and the frequency domain is discretized according to the minimal resolvable frequency of the discrete Fourier transform. To fix the notation, let us label each time segment as a and each frequency bin as k ; then, each pixel (a, k) within the time-frequency plane gets assigned a binary value given by

$$n_a[k] = \begin{cases} 1, & \rho_a[k] > \rho_{th} \\ 0, & \rho_a[k] < \rho_{th} \end{cases}, \quad (3.5)$$

where $\rho_a[k]$ refers to the normalized power within the k -th frequency component of the a -th data segment.

Figure 3.1 shows qualitatively the construction of a binary map for a continuous monochromatic signal under the presence of noise. Such a signal could be caused by a continuous wave source located at one of the equatorial poles, where the Doppler shift is practically negligible. Even though the monochromatic signal is recovered as a horizontal strip along the binary map, certain power bins produced by noise fluctuations also survive the thresholding procedure. We will expose how the *SkyHough* pipeline automatically deals with such undesired artifacts during the forthcoming explanation of the incoherent step of the algorithm.

3.1.2 Incoherent Step

The strategy is to rank the parameter space templates according to its prominence as normalized power tracks in the data. Each template $\vec{\lambda}$ will get assigned a *number count* $n(\vec{\lambda})$ by *counting* the amount of bins above the power threshold in the binary map. More formally, let us define $T(\vec{\lambda})$ as the frequency

evolution track yielded by the template $\vec{\lambda}$ onto the binary map; then, the number count $n(\vec{\lambda})$ is defined as

$$n(\vec{\lambda}) = \sum_{(a,k) \in T(\vec{\lambda})} n_a[k]. \quad (3.6)$$

We can use figure 3.1 to illustrate this ranking procedure: For instance, if we were to take the template corresponding to the monochromatic signal at frequency $6 \cdot (\delta f)$, we would end up with a number count of 11, since the whole horizontal track is filled with power above the threshold; on the other hand, a spinning-down signal starting at the top left corner of the binary map and ending at the bottom right corner would get assigned a number count of 2.

This description of the number count treats every time-frequency bin in the same manner; however, as discussed in chapters 1 and 2, neither the detector's response is isotropic, nor its noise is white. Thus, it is natural to ask for a different definition of the number count, in order to properly treat both features.

To do so, a *weighted* number count was introduced [47]; that is, rather than equally contributing to the total count, each bin (a, k) within a track $T(\vec{\lambda})$ will be weighted by a quantity $w_a[k]$, which will take into account the noise floor of the detector and the amplitude modulation suffered by the signal due to the antenna pattern functions

$$n(\vec{\lambda}) = \sum_{(a,k) \in T(\vec{\lambda})} w_a[k] n_a[k]. \quad (3.7)$$

The choice of $w_a[k]$ was discussed in [60] and [47]. The argument describes the sensitivity of the search in terms of the spectral power of a signal measured by an interferometric detector (1.10), averaging over parameters which directly affect its amplitude. Detector noise enters as the power spectral density S_n ; antenna pattern functions $F_{+,\times}$ enter as polarization averaged functions $\langle F_{+,\times}^2(\vec{n}; t) \rangle_\psi$. For a power bin (a, k) , the corresponding weight is given by

$$w_a[k] = \mathcal{N} \frac{\langle F_{+a}^2[k] \rangle_\psi + \langle F_{\times a}^2[k] \rangle_\psi}{S_{na}[k]}, \quad (3.8)$$

where \mathcal{N} represents an arbitrary normalization constant and the notation $f_a[k]$ refers to the evaluation of f using the parameters which correspond to the k -th frequency bin of the a -th data segment. Reading its definition, it is clear that $w_a[k]$ increases the contribution of those bins with low floor noise and good orientation with respect to the interferometric detector.

The statistical significance of this results has to be further discussed, since any measurement is poisoned by different kinds of noise, which may trigger power bins above the first threshold, or even build a significant number count, as certain spectral disturbances do. We will devote section 3.3 to do so, describing the number count statistics, as well as the optimal choice of thresholds to discriminate among noise and signal.

3.2 Implementation

After the description of the main algorithm, we need to tackle two issues in order to describe the whole picture. The first one, here discussed, deals with the actual implementation of the method from a computational point of view; the second one, with its own section, will deal with the appropriate choice of thresholds from a statistical point of view.

We will describe the operational choices that were made to implement the SkyHough pipeline as described by [46]. The actual implementation can be found within the LIGO Algorithm Library (LAL) [73].

3.2.1 Length of Coherent Segments

The first issue to address concerns the time length choice of the coherently analyzed data segments T_{coh} , described in section 3.1.1. We recall that this parameter effectively defines the grid step of the time-frequency plane, since the resolution of both subspaces fulfills an inversely proportional relationship; hence, its choice will directly affect the way we *see* signals in the binary maps.

The main argument goes as follows: The frequency evolution of a signal will be mainly affected by two modulations, namely the natural spin-down of the source and the Doppler shift due to Earth's movement in the SSB frame. If such modulations spread the main components of the signal across several frequency bins, the thresholding procedure will not be able to distinguish its significance against the noise; hence, we will choose a coherent time length T_{coh} such that the frequency modulations are kept within a single frequency bin at each data segment. We impose this condition constraining the maximum frequency shift to one *half* of a frequency bin.

Let us abstract the maximum frequency shift as $|\dot{f}|_{\text{max}}$, whichever its origin and sign (Doppler modulation, Spin-down/up of the source, ...). The aforementioned constraint can be expressed as

$$|\dot{f}|_{\text{max}} T_{\text{coh}} < \frac{1}{2} \delta f, \quad (3.9)$$

which directly translates to

$$T_{\text{coh}} < \sqrt{\frac{1}{2|\dot{f}|_{\text{max}}}}. \quad (3.10)$$

We will consider the Doppler shift as the main source of frequency modulation¹. Start by differentiating in equation (2.13) with a fixed $\hat{f}(t) = \hat{f}$ (we neglect the spin-down effect against the Doppler Shift):

$$\frac{df}{dt} \simeq \frac{\hat{f}}{c} \frac{d\vec{v}}{dt} \cdot \vec{n} \leq \frac{\hat{f}}{c} \left| \frac{d\vec{v}}{dt} \right| \equiv |\dot{f}|_{\text{max}}. \quad (3.11)$$

The main contribution to the Doppler shift comes from the daily rotation of Earth around its own axis, which means that we have to take into account the centripetal acceleration of a detector onto the surface of Earth as

$$\left| \frac{d\vec{v}}{dt} \right| = \frac{v_{\oplus}^2}{R_{\oplus}} = \frac{4\pi^2 R_{\oplus}}{T_{\oplus}^2}, \quad (3.12)$$

where v_{\oplus} stands for the magnitude of the velocity of the surface of Earth around its axis, T_{\oplus} is the duration of a day and R_{\oplus} the radius of Earth.

Substituting numerical values in (3.12) and arranging (3.11) into (3.10) to get a meaningful constraint, we obtain

$$T_{\text{coh}} < 50\text{min} \times \sqrt{\frac{500\text{Hz}}{\hat{f}}}. \quad (3.13)$$

A typically used value is $T_{\text{coh}} = 1800\text{s}$, although one may combine several coherent time lengths to cover different regions of the parameter space.

¹This fact will have consequences on the allowed spin-down parameters to search for $\{f_s\}$, as we will briefly discuss in section 3.2.3

3.2.2 Discretization of the Parameter Space

The computational implementation of an algorithm curses numbers to live in a finitely-resolved, discrete, space. This forces us to carefully review the parameter space introduced in by (3.1), since the difference between two adjacent templates will not be arbitrarily small; rather, there will exist a resolution step

$$\delta\vec{\lambda} \equiv (\delta f, \{\delta f_s\}, \delta\theta), \quad (3.14)$$

where $\delta f = T_{\text{coh}}^{-1}$ as already stated, δf_s represents the resolution of the spin-down parameter's subspace and $\delta\theta$ represents the resolution onto the celestial sphere, which is taken equal for both degrees of freedom of \mathbb{S}^2 .

The resolution of spin-down parameters can be easily derived from an operational point of view. Since they enter into the game as coefficients of a Taylor expansion, we will take the resolution of the s -th spin-down coefficient δf_s as the minimum value that would produce a minimal distinguishable change in frequency δf by itself during the whole time of measurement T_{obs} ; that is,

$$\delta f = \frac{T_{\text{obs}}^s}{s!} \delta f_s \rightarrow \delta f_s = \frac{s!}{T_{\text{obs}}^s} \delta f. \quad (3.15)$$

As for the sky resolution, we need to understand the way in which the sky position of a source affects its frequency modulation. Take the Doppler shift as expressed in (2.13). At a certain time t , the locus of sky positions $\vec{n} \in \mathbb{S}^2$ consistent with a certain frequency observation $f(t)$ is given by

$$\cos \theta = \frac{\vec{v} \cdot \vec{n}}{v} = \frac{c}{v} \frac{f(t) - \hat{f}(t)}{\hat{f}(t)}, \quad (3.16)$$

where θ represents the angle between the sky position of the source \vec{n} and the velocity vector of the detector \vec{v} , and the lack of arrow denotes the modulus of the vector. This simple equation reveals the actual structure of the Doppler shift: Sky positions consistent with a certain frequency f are located onto *rings* of the celestial sphere with its center located at the tip of the velocity vector of the detector. However, the finite frequency resolution *dilates* rings into annuli, since all the frequencies within a frequency bin δf are effectively indistinguishable by the algorithm. The thickness of such annuli can be easily derived from (3.16), yielding

$$\delta\theta \simeq \frac{c}{v} \frac{\delta f}{\hat{f} \sin \theta}. \quad (3.17)$$

It is clear that the thickness of the annuli reaches a minimum at $\theta = \frac{\pi}{2}$, and grows monotonically as we get parallel to the velocity vector; hence, we will choose a fraction of the minimum thickness to tessellate the celestial sphere

$$\delta\theta \equiv \frac{1}{p} \frac{c}{v} \frac{\delta f}{\hat{f}}, \quad (3.18)$$

where the *pixel factor* p can be tuned to increase the sky resolution according to the available computational capabilities. This resolution will be used to set up a grid of *sky patches* onto the sphere. The relation between annuli and sky patches can be set up through the sky patch center: A sky patch belongs to a particular annulus if its center is located into such annulus. This relation ensures a well representation of the Doppler shift on the discretized version of the sphere.

3.2.3 Spin Down ranges

The amount and range of spin-down parameters $\{f_s\}$ need to be considered in order to understand what physics are representable under a certain choice of such. It will become clear that the main physical

concern when constraining the spin-down parameters is related to the *age* of the source; that is, the amount of time the spin-down could modulate the source frequency.

We start by restricting the amount of spin-down parameters to be considered on our searches: The same argument that led us to define the spin-down resolution can serve us to put a constraint on the maximum value allowed for a spin-down parameter f_s . If we accept the description of the frequency evolution as a Taylor series, the *maximum* allowed value for the s -th spin-down parameter is given in terms of the age of the source τ as

$$f_s^{\max} = s! \frac{\hat{f}_{\max}}{\tau^s}. \quad (3.19)$$

Source's age τ is defined as the amount of time required by the source to achieve the current frequency f_0 assuming a continuous energy emission solely by gravitational waves.

Now, typical astrophysical continuous wave sources are older than the affordable observation times $T_{\text{obs}} \ll \tau$. This implies that the *maximum* allowed spin-down value decreases *faster* than the spin-down resolution (3.15) as s increases; hence, there will exist a maximum number s_{\max} for which spin-down parameters are resolvable, establishing a meaningful way to limit the number of spin-down parameters in terms of the age of the targeted sources. Equivalently, one could *fix* the value s_{\max} according to the affordable computational capabilities, obtaining a minimal detectable source age in exchange.

Regarding previous considerations on the main source of modulation during section 3.2.1, we can easily relate the constraint of less than half bin of modulation to the s -th spin-down parameter as

$$f_s^{\max} T_{\text{coh}}^s < s! \frac{\delta f}{2} \quad (3.20)$$

which, in combination with (3.19), yields the corresponding constraint on the age of the source

$$\tau > \left(\frac{2\hat{f}_{\max} T_{\text{coh}}^{s+1}}{s!} \right)^{1/s}. \quad (3.21)$$

Summarizing, the number of spin-down values is related to the age of the sources through its resolution, meaning that one can choose either to fix a source age in order to target a particular kind of physics, or to constraint the amount of parameters to construct an affordable, wide-spread search. Alternatively, one could choose to disregard the spin-down contribution in favor of the Doppler shift, automatically obtaining the corresponding constraint over the age of the source (3.21).

The *SkyHough* pipeline takes the latter approach with only *one* spin down parameter, implying

$$\tau > 103\text{yr} \times \frac{\hat{f}_{\max}}{1000\text{Hz}} \left(\frac{T_{\text{coh}}}{1800\text{s}} \right)^2. \quad (3.22)$$

As discussed in [80], this choice allows for a certain range of higher order spin-down values to be taken into account, namely those which would not shift the frequency of the signal more than half of a frequency bin during a coherent time T_{coh} .

3.2.4 Partial Hough Maps

We will devote a last section to comment on the actual way the Hough Transform is implemented within the *SkyHough* pipeline.

It was already discussed during section 3.1.2 that each template $\vec{\lambda}$ can be related to a certain power track on the time-frequency plane $T(\vec{\lambda})$. Although one could run blindly along the parameter space to compute

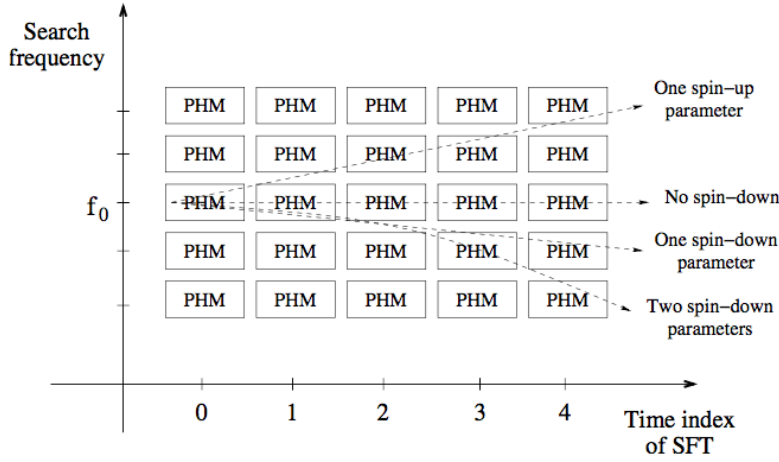


Figure 3.2: The driving effect on the frequency evolution of the signal is given by the spin-down parameters, which enter into the frequency modulation as polynomial terms: No spin-down implies a constant frequency; one spin-down parameter implies a linear evolution; two spin-down parameters imply a parabolic evolution. Image Source: [46]

the number count of each template, the sky annuli disposition given by (3.16) suggests a more efficient way to do so.

Consider a certain frequency to look for f_0 at a certain time stamp a . We already showed that the compatible sky positions are located at thick rings onto the celestial sphere; however, one could argue that this structure is just a histogram on the (α, δ) plane of equatorial sky positions, where the compatible bins are suitably labeled; let us refer to this structure as *Partial Hough Map* (PHM).

Hence, the number count computation can be automatically done for each sky patch, since everything one has to do is to compute and add the corresponding PHMs for different time stamps according to the frequency evolution of the signal, obtaining the *Total Hough Map* (THM) for each set of values of frequency and spin-down parameters $(f_0, \{f_s\})$. Figure 3.2 shows the path of PHMs to be summed, depending on the spin-down parameters taken into account.

Partial Hough Maps can be re-used by several frequency bins, since the sky annuli structure is highly insensible to frequency changes thanks to the monotonicity of its thickness. A formal explanation of this fact can be found in [46]. Therefore, instead of recomputing the set of PHM for each set of $(f_0, \{f_s\})$, we will store them inside *look up tables* (LUT), lowering the computational load of the search.

This reduces the number count computation to a sum of *re-used* PHM for each set of $(f_0, \{f_s\})$. As a result, the histogram on the parameter space will be described as a set of Total Hough Maps.

Other implementations of the Hough method differ in this step, performing the construction over frequency and spin-down for a fixed sky position [13].

3.3 Statistics

We are left with the suitable choice of ρ_{th} and a discussion on the statistical significance of the number count (3.7). We will introduce the appropriate statistical machinery to distinguish between actual signals and spurious noise fluctuations: First, we will study the way in which noise will enter our searches via the statistical properties of the number count; then, we will use such information to construct the appropriate

thresholds, which will be able to state the significance of a certain template in terms of probabilities.

3.3.1 Normalized Power distribution

Consider an additive, *stationary, zero-mean*, white Gaussian noise $n[j]$. Following the exposed algorithm, we can express the noise normalized power as

$$2\rho[k] = z_{\mathcal{R}}[k]^2 + z_{\mathcal{I}}[k]^2 \quad (3.23)$$

where

$$z_{\mathcal{R}}[k] = \frac{\sqrt{2}\mathcal{R}(\tilde{x}[k])}{\sqrt{\langle |\tilde{n}[k]|^2 \rangle}}, \quad z_{\mathcal{I}}[k] = \frac{\sqrt{2}\mathcal{I}(\tilde{x}[k])}{\sqrt{\langle |\tilde{n}[k]|^2 \rangle}}, \quad (3.24)$$

$\mathcal{R}(\cdot)$ and $\mathcal{I}(\cdot)$ represent the real and imaginary part functions, respectively, and the notation has been kept consistent with the one used during section 3.1.

Being $n[j]$ Gaussian, it is clear that $\mathcal{R}[\tilde{n}[k]]$ and $\mathcal{I}[\tilde{n}[k]]$ are also Gaussian, since they can be described as a convolution of Gaussian random variables. Moreover, its whiteness justifies the fact that both of them have the same variance, which is given by $\langle |\tilde{n}[k]|^2 \rangle / 2$; hence, $z_{\mathcal{R}}[k]$ and $z_{\mathcal{I}}[k]$ are both Gaussian random variables with unit variance.

This implies that $2\rho_k$ is a sum of *two* squared Gaussian variables, i.e. a χ^2 distributed variable with *two* degrees of freedom and non-centrality parameter given by

$$\lambda[k] = \langle z_{\mathcal{R}}[k] \rangle^2 + \langle z_{\mathcal{I}}[k] \rangle^2 = \frac{4 |\tilde{h}[k]|^2}{T_{\text{coh}} S_n[k]}. \quad (3.25)$$

The non-centrality parameter describes the significance of a signal proportional to its normalized power, raising the mean power value over the zero mean noise. Thus, the distribution of $\rho[k]$ under the presence of a signal described by a non-centrality parameter $\lambda[k]$ is given by

$$\begin{aligned} p(\rho[k]|\lambda[k]) &= 2\chi^2(2\rho[k]|2, \lambda[k]) \\ &= \exp\left(-\rho[k] - \frac{\lambda[k]}{2}\right) I_0\left(\sqrt{2\lambda[k]\rho[k]}\right), \end{aligned} \quad (3.26)$$

where $I_0(\cdot)$ is the modified Bessel function of zeroth order. As the signal fades into the noise, the distribution approaches an exponential shape. Hereafter we shall drop the discrete index $[k]$ in order to keep a clear notation, understanding that ρ and λ refer to one of the Fourier frequency bins.

Given a threshold ρ_{th} to construct binary maps, the probability of selecting a frequency bin containing a signal with non-centrality parameter λ is given by

$$\eta(\rho_{\text{th}}|\lambda) = \int_{\rho_{\text{th}}}^{\infty} p(\rho|\lambda) d\rho, \quad (3.27)$$

i.e. the probability of ρ triggering above the specified threshold due to the effect of an underlying signal.

Alas, one needs special care when a hard threshold is set: It could happen (and, certainly, it *will* happen) that a spurious noise fluctuation triggers *above* the threshold; even worse, we could be completely missing a weak signal whose power is *not* enough to surpass the established threshold.

The former effect will be treated in terms of the *false alarm* probability

$$\alpha(\rho_{\text{th}}) = \int_{\rho_{\text{th}}}^{\infty} p(\rho|0) d\rho = e^{-\rho_{\text{th}}}, \quad (3.28)$$

which computes the probability of selecting a power bin made out of pure noise; the latter effect will be represented by the *false dismissal* probability

$$\beta(\rho_{\text{th}}|\lambda) = 1 - \eta(\rho_{\text{th}}|\lambda) = \int_0^{\rho_{\text{th}}} p(\rho|\lambda) d\rho, \quad (3.29)$$

which encloses the probability of missing a frequency bin containing a signal.

We will close this set of definitions by connecting the peak selection probability η with the false alarm probability α . First, it is clear by definition that $\eta(\rho_{\text{th}}|0) = \alpha(\rho_{\text{th}})$, since the peak selection probability in the absence of signal relies entirely on the noise. Moreover, continuous wave signals tend to be weaker than the surrounding noise, allowing for a *weak signal* approximation [19]; that is, the spectral intensity of a signal is small against the surrounding noise power. In our terms, $\lambda \ll 1$, and η can be easily given in a closed form as

$$\eta(\rho_{\text{th}}|\lambda) = \alpha(\rho_{\text{th}}) \left\{ 1 + \frac{\rho_{\text{th}}}{2} \lambda + \mathcal{O}(\lambda^2) \right\}. \quad (3.30)$$

3.3.2 Number Count Distribution

The number count statistic can be easily derived once normalized power is understood. We will derive the probability distribution of measuring a number count $n(\vec{\lambda})$ in a particular pixel $\vec{\lambda}$ of the Total Hough Map. We recall that the number count of a template is constructed as a sum of selected normalized power bins along the time-frequency path yielded by such template, as expressed in (3.6).

Consider a binary map constructed using N coherent segments. The probability of obtaining a number count n out of a certain time-frequency path can be described using a set of N Bernoulli trials: Each of the coherent segments will yield an increase in the number count *if* its frequency bin surpasses the specified threshold, being η the probability of doing so. Thus, the probability distribution of measuring a number count n in a particular pixel for a certain signal intensity λ is given by the binomial distribution

$$p(n|\rho_{\text{th}}, \lambda) = \binom{N}{n} \eta^n (1 - \eta)^{N-n}. \quad (3.31)$$

Let us define $\vec{w} = (w_1, \dots, w_N)$ as the set of weights involved in the number count computation. The mean value μ and variance σ^2 are given by

$$\mu = A\eta, \quad \sigma^2 = \|\vec{w}\|^2 \eta(1 - \eta), \quad (3.32)$$

where

$$A = \sum_{i=1}^N w_i, \quad \|\vec{w}\|^2 = \sum_{i=1}^N w_i^2. \quad (3.33)$$

For the case of unitary equal weights $w_i = 1, \forall i$, both quantities reduced to N . In the absence of a signal, we recover the same expressions with $\eta = \alpha$. This latter case will be referred to as the *noise statistics*.

Although we have not given actual values to A , $\|\vec{w}\|^2$ and α yet, it will become clear during the following section that their typical ranges allow for the use of the Central Limit Theorem, describing the binomial distribution using a Gaussian approximation

$$p(n|\rho_{\text{th}}, \lambda) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(n-\mu)^2}{2\sigma^2}}, \quad (3.34)$$

where the mean and standard deviation are consistent with those of the binomial distribution 3.32.

The final step of the algorithm sets a threshold n_{th} on the number count. As before, such a choice will breed a false alarm and false dismissal probabilities in a completely analogous way as that of the previous section. Thanks to the Gaussian approximation, we are able to quote both the false alarm α_H and false dismissal β_H probabilities as

$$\alpha_H(\rho_{\text{th}}, N, n_{\text{th}}) = \int_{n_{\text{th}}}^{\infty} p(n|\rho_{\text{th}}, 0)dn = \frac{1}{2}\text{erfc}\left(\frac{n_{\text{th}} - N\alpha}{\sqrt{2N\alpha(1-\alpha)}}\right), \quad (3.35)$$

$$\beta_H(\rho_{\text{th}}, N, n_{\text{th}}, \lambda) = \int_0^{n_{\text{th}}} p(n|\rho_{\text{th}}, \lambda)dn = \frac{1}{2}\text{erfc}\left(\frac{N\eta - n_{\text{th}}}{\sqrt{2N\eta(1-\eta)}}\right), \quad (3.36)$$

where erfc represents the complementary error function and the expressions for the mean and standard deviation have been made explicit in terms of the peak selection probabilities.

Let us explicitly distinguish between α, β and α_H, β_H : The former are related to the selection of normalized power bins according to the power threshold ρ_{th} ; the latter are related to the selection of templates in the parameter space, according to the number count threshold n_{th} .

3.3.3 Optimal Choice of thresholds

Two ways can be taken in order to set appropriate values to the aforementioned thresholds ρ_{th} and n_{th} : Either rely on a well-known statistical result, the *Neyman-Pearson* criterion, or explicitly maximize the effectiveness of the search by a suitable choice of ρ_{th} . We will use the latter for convenience, although both methods yield completely equivalent results within the weak signal regime.

So far, we have implicitly referred our statistical measurements, such as the number count n , to an arbitrary null value; that is, neither we know how significant is a certain number count, nor we are able to compare it to another value *solely by itself*. This lack of information can be emended with a new variable, the critical ratio

$$\Psi(n) = \frac{n - \mu}{\sigma} = \frac{n - A\alpha}{\sqrt{\|w\|^2 \alpha(1-\alpha)}}, \quad (3.37)$$

where μ and σ are taken to be the ones corresponding to *noise*. Essentially, it computes how far is a number count n with respect to the average noise in units of its standard deviation.

Given that higher critical ratios imply more significant candidates, we will choose a threshold ρ_{th} which maximizes the typical critical ratio. That is, take its average value, according to the number count statistic

$$\langle \Psi \rangle = \frac{N\eta - N\alpha}{\sqrt{N\alpha(1-\alpha)}}, \quad (3.38)$$

and maximize with respect to ρ_{th} , taking into account the linear expansion of the peak selection probability (3.30) and the definition of false alarm peak selection probability (3.28):

$$\frac{\partial \langle \Psi \rangle}{\partial \rho_{\text{th}}} = 0 \rightarrow \log \alpha = 2(\alpha - 1). \quad (3.39)$$

This transcendental equation can be numerically solved to yield $\rho_{\text{th}} \simeq 1.6$ or $\alpha \simeq 0.2$. It is worth mentioning the independence of this result on the normalization of the weights given in (3.8): Any transformation $\omega_i \rightarrow \mathcal{N}\omega_i, \forall i$ will be factored out from (3.37); hence, we can freely choose the value of A , according to our needs. For instance, $A = N$ yields $n \leq N$.

As for the number count threshold n_{th} , it can be set by establishing a false alarm probability α_H and inverting (3.35). This choice has to be carefully made: A low false alarm rises the threshold, diminishing the number of candidates to be followed up; in exchange, false dismissal probability increases, potentially punishing the sensitivity of the search. The actual line of reasoning behind the theoretical sensitivity of a search, which goes beyond the scope of the present work, can be found in [46]. Further sections will expose the actual method we used to quantize the sensitivity of a search.

3.4 Post Processing

The first result of a search in a region of the parameter space is given in terms of a *toplist*, which stores the best candidates in terms of its critical ratio. One could either set a minimal threshold to enter in the toplist or specify a certain amount of candidates to be stored in order to fulfill computational requirements.

Anyway, the result is a cloud of significant points in the parameter space, the origin of which could be either a noise disturbance or an actual astrophysical signal. We proceed to describe the post processing steps used to further discriminate both candidates. There are two steps involved:

The first one explodes the existence of multiple interferometric detectors with similar sensibilities. If a continuous gravitational wave signal reaches the Earth, the response of every detector has to be consistent; in another way, the parameters corresponding to a particular signal have to be present in every considered toplist. Such a coincidence has to be done using a *common* epoch; that is, source parameters yielded by two detectors, have to be taken at the same GPS reference time. To do so, frequencies are shifted using its corresponding spin-down and the shift between reference times.

Given that parameter space is discretized, a continuous wave signal can be triggered by several templates, since none of them actually represents its *exact* parameters. Hence, the second step intends to cluster surviving templates from the previous step into a single region of the parameter space.

Clusters due to a signal are expected to trigger more than one candidate of the parameter space; hence, any single-candidate cluster can be disregarded as a spurious noise fluctuation. This imposition is referred to as *population veto*.

Alternatively, one could recompute the significance of a cluster using other, more significant statistics. For instance, O2 analysis used the \mathcal{F} -statistic to discard noisy clusters, re-computing its significance with different coherent lengths.

3.4.1 Parameter Space Distance

A distance function is introduced to construct template coincidences in the parameter space. Let $\vec{\lambda}^a$ and $\vec{\lambda}^b$ be two templates in the parameter space (3.1). The *SkyHough* distance is defined as

$$d_H(\vec{\lambda}^a, \vec{\lambda}^b) = \sqrt{\left(\frac{\Delta f}{\delta f}\right)^2 + \left(\frac{\Delta f_1}{\delta f_1}\right)^2 + \left(\frac{\Delta \theta}{\delta \theta}\right)^2}, \quad (3.40)$$

where Δf represents the distance in the frequency subspace, Δf_1 represents the distance in the spin-down subspace and $\Delta \theta$ represents the distance in the sky subspace. Each contribution is expressed in terms of bins in the corresponding subspace δf , δf_1 , $\delta \theta$, the expression of which was shown in section 3.2.

The frequency and spin-down contributions are taken as the absolute difference between both templates

$$\Delta f = |f_0^a - f_0^b|, \quad \Delta f_1 = |f_1^a - f_1^b|. \quad (3.41)$$

As for the sky contribution $\Delta\theta$, two prescriptions have been proposed so far.

During the O1 data analysis [8], $\Delta\theta$ was taken as the *geodesic* distance over the celestial sphere between the two candidates

$$\Delta\theta = \arctan \frac{|\vec{n}^a \times \vec{n}^b|}{\vec{n}^a \cdot \vec{n}^b}. \quad (3.42)$$

The drawback comes from the structure of the Doppler modulation of the signal. As already discussed, the Doppler shift is the main source of sky location of a continuous wave signal. According to the annuli discussion in section 3.2.2, the loci formed by sky locations *parallel* to the velocity vector of the detector is composed by the thickest annuli, spreading candidates along wider regions. This region corresponds to the vicinity of the ecliptic plane.

The proposed solution during the O2 data analysis [9] was another prescription of $\Delta\theta$. Instead of a geodesic distance over the sphere, candidates were *projected* to the ecliptic plane, taking $\Delta\theta$ as the euclidean distance onto such plane.

For an arbitrary sky position expressed in equatorial coordinates $\vec{n}(\alpha, \delta)$, the corresponding cartesian coordinates onto the celestial sphere (ξ, η, ζ) are defined as

$$\begin{cases} \xi = \cos \delta \cos \alpha \\ \eta = \cos \delta \sin \alpha \\ \zeta = \sin \delta \end{cases}. \quad (3.43)$$

The conversion to cartesian coordinates with respect to the ecliptic plane is given by a rotation around the ξ axis according the obliquity of the ecliptic φ

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \varphi & \sin \varphi \\ 0 & -\sin \varphi & \cos \varphi \end{pmatrix} \begin{pmatrix} \xi \\ \eta \\ \zeta \end{pmatrix}. \quad (3.44)$$

Then, the sky contribution to the parameter space distance between two candidates is given by

$$\Delta\theta = \sqrt{(x^a - x^b)^2 + (y^a - y^b)^2}. \quad (3.45)$$

This projection lets the region of the poles almost untouched; however, signals near the ecliptic plane are located *closer*, since the main source of delocalization, given by the z component in (3.44), is completely disregarded.

3.4.2 Coincidence and Clustering

The first step of the post processing uses two toplist coming from two different detectors. During the last scientific runs, those where the L1 and H1 LIGO interferometers located at Livingston and Hanford respectively. Each candidate in one of the toplit, say $\vec{\lambda}^H$, was compared to every candidate in the other toplist $\vec{\lambda}^L$. If the distance between a pair of candidates fulfilled $d_H(\vec{\lambda}^H, \vec{\lambda}^L) \leq r_c$, a new candidate was computed as the weighted average of $\vec{\lambda}^L$ and $\vec{\lambda}^H$ according to its critical ratio. The result of this step is a new toplist of coincident candidates among both detectors. We recall that *any* pair of candidates can produce a coincidence; that is, a candidate from one toplist could produce a coincidence with every single candidate in the other toplist.

Let us define $\Lambda^c = \{\vec{\lambda}_1^c, \dots, \vec{\lambda}_M^c\}$ as the obtained toplist of coincident candidates. The second step intends to cluster them into bigger sets. To do so, candidates are selected pairwise; for any pair $i, j \in [1, M], i \neq j$

we compute its parameter space distance. If they are located closer than a certain window r_{cl} , i.e. $d_H(\vec{\lambda}_i^c, \vec{\lambda}_j^c) \leq r_{cl}$, both candidates are included into the same cluster in a transitive fashion.

So far, the suitable choice of r_c and r_{cl} has been manually calibrated for each analysis. During the O1 analysis, using the geodesic sky distance (3.42), the choice was $r_c = r_{cl} = \sqrt{14}$ in units of parameter space discrete bins. For the O2 analysis, where the projected sky distance was introduced (3.45), the choice was $r_c = 3$ and $r_{cl} = \sqrt{14}$. The use of a lower r_c reduces the chance of random fluctuations to coincide, while a higher r_{cl} groups candidates from a common origin.

In the end, this procedure lows the false alarm probability of a search by using extensively the properties of a gravitational waves measurement. Any surviving cluster becomes a region to be followed up by deeper methods in order to refine its main properties.

The sixth LIGO scientific run and its Continuous Wave Mock Data Challenge

We expose the Mock Data Challenge performed during the sixth LIGO scientific run by the continuous wave pipelines. In addition to being a clear comparison between different search strategies, it will be useful to compare the behavior of further improvements into the *SkyHough* pipeline.

4.1 LIGO S6 Data

The sixth LIGO scientific run (S6) [3] was the last measurement period of the LIGO detectors before its major upgrades to the current Advanced LIGO detectors. Its spanned along fifteen months, from 8th July 2009 (GPS 931035615) to 17th October 2010 (GPS 971622015), although the useful time is much shorter due to commissioning periods undergoing in the detectors. The dataset is freely available under the Gravitational Wave Open Science Center (GWOSC) [79], maintained by the LIGO and Virgo scientific collaboration.

The run was divided in four segments, labeled as A-D, according to changes in the detector performance due to multiple reasons. Segment A run approximately for two months, stopping for a commission break. As for segment B, it ended the rest of the year, ending with a commission break as well. Regarding epochs C and D, they spanned a continuous nine month period of operation, being the division marked by the inclusion of the Virgo detector into the search.

Data quality, understood as a measure of instrumental stability of the detector, was measured in terms of the *duty factor*, the fraction of the total run time which is useful for scientific purposes. Each continuous

Segment	Median duration (mins)	Longest duration (hours)	Total live time (days)	Duty factor (%)
S6A	54.0/39.3	13.4/11.8	27.5/25.6	49.1/45.7
S6B	75.2/17.3	19.0/21.3	59.2/40.0	54.3/38.0
S6C	82.0/67.5	17.0/21.4	82.8/82.3	51.4/51.1
S6D	123.4/58.2	35.2/32.6	74.7/75.2	63.9/64.3

Table 4.1: Characteristics of the S6 run for the H1/L1 LIGO interferometers, located at the LIGO Hanford Observatory and LIGO Livingston Observatory, respectively.




Figure 4.1: Amplitude Spectral Density of the LIGO detectors during S6. The black line shows the predicted design sensitivity for the instruments [21], while the green and red lines show the actual, measured ASD. Source: [3].

period of operation under nominal detector conditions is called *science segment*. Its duration is submitted to the detector, ending when the noise levels of the instrument surpass the electronically controlled region. Alternatively, science segments can be manually ended to perform maintenance tasks.

Table 4.1 summarizes the results delivered by both LIGO detectors after the scientific run. The overall tendency points to an increase in the duty factor for both detectors as the run progressed. The same tendency can be observed on the duration of the science runs. The short duration science segments during the early days are mostly due to the significantly poor detector stability suffered by L1, especially during the B segment of the run. Increasing instrumental stability shows the deeper understanding achieved by the LIGO collaboration regarding noise sources of the detector, as well as an improvement in the control systems used to maintain the optimal working point of the instrument.

As for the sensitivity to gravitational waves, the usual measurement is the strain amplitude spectral density of detector output, which we called simply amplitude spectral density (ASD), shown in figure 4.1. As already discussed during section chapter 1, the lower region of the spectra is due to seismic noise and its coupling to the mechanical suspensions of the mirror; the substantial gap among L1 and H1 can be related to the prototype isolation installed in the former. Higher region of the spectra can be related to thermal noise or quantum fluctuations of the laser's properties; however, there were some structures the origin of which was never fully understood.

Regarding continuous wave searches, the most critical feature is related to the existence of narrow-band *line* structures into the spectra, which correspond to persistent noise disturbances with a specific frequency. Many of such disturbances are a direct consequence of the design and operation of the interferometers: The 60Hz line is due to the alternating current power line used in the U.S.; instrumental suspensions produce a prominent line around 350Hz; calibration lines are located on purpose to study the response function of the detector; harmonics of noise lines are also into play, such as the 120Hz power line harmonic, clearly seen in figure 4.1. There were another sets of undesired lines, the origin of which was poorly understood; some of them poisoned both instruments, some of them just one. In particular, two sets of 2Hz and 16Hz harmonics appeared coherently in *both* observatories, disturbing the searches for long duration gravitational wave signals.

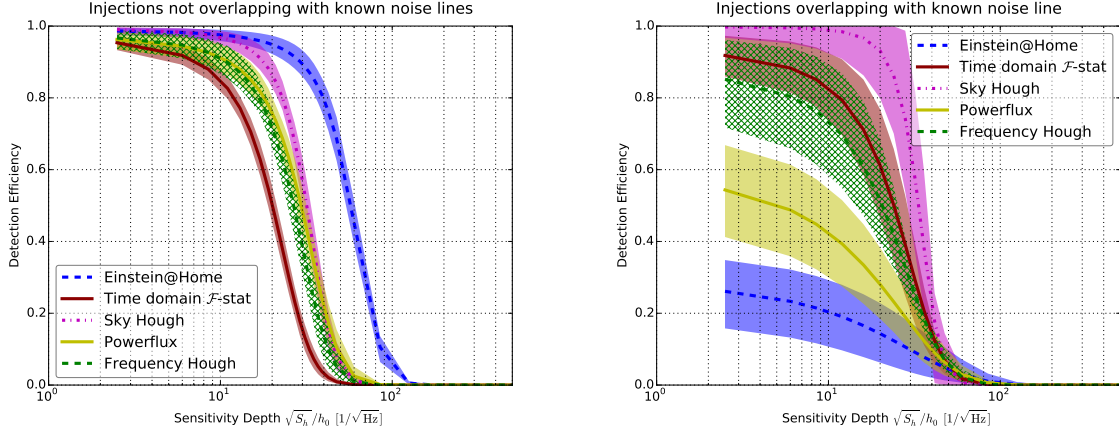


Figure 4.2: Efficiency curves for the Mock Data Challenge performed over the LIGO S6 science run. Different pipelines are represented by different colors. The shaded region around each line represents the statistical uncertainty on the detection efficiency measured for the search implementation and the data used in the MDC. **Left:** Efficiency curve using injections which did not overlap with a known noise line. **Right:** Efficiency curve using injections which did overlap with a known noise line. Source: [80].

Continuous wave pipelines chose to analyze the final seven months of the S6 data set due to the spectral contamination during the initial periods of the run [77]. Moreover, frequency bins which were clearly contaminated by noise disturbances were removed from the analysis. In all, the S6 analysis was able to yield better results than the previous S5 analysis, even though an explanation for the increasing number of lines remained pending.

4.2 The S6 Mock Data Challenge

Mock Data Challenges (MDCs) are a way to test the performance of a pipeline using real (or close to real) conditions. In this case [80], a set of artificial continuous gravitational wave signals were generated and introduced (*injected*) into the S6 data set. The challenge was to find such signals using the different pipelines involved in continuous wave searches. After the challenge, one is able to understand the performance of each pipeline with respect to others regarding sensitivity, robustness against noise disturbances and computational costs.

3110 signals consistent with an isolated, spinning neutron star were injected into real S6 data, from GPS 931035615 to GPS 971622272. Injections were placed in frequency bands at 0.5Hz intervals. Sky position was taken as a random variable isotropically distributed over the celestial sphere, while spin down was drawn from a uniform distribution in log space within $[-1 \times 10^{-9}, -1 \times 10^{-18}]$ Hz/s for a 95% of the signals and $[1 \times 10^{-18}, 1 \times 10^{-13}]$ Hz/s for the remaining 5%. A 25% of the signals got assigned a braking index $n \in [5, 7]$, which was reflected in non-null second and third frequency derivatives. Such brake indexes can be related to neutron stars with gravitational waves as the main emission channel. The amplitude of the signals h_0 was described in terms of the signal-to-noise ration of a coherent search, drawing its values from a uniform distribution. The effective noise amplitude spectral density $\sqrt{S_n}$ was computed using an inverse-variance-weighted harmonic average between both detectors. As for the nuisance parameters ψ, ϕ_0 and $\cos \iota$, they were randomly drawn from uniform distribution with ranges $[-\pi/4, \pi/4]$, $[0, 2\pi]$ and $[-1, 1]$, respectively.

The signals were generated and injected into S6 data using `lalapps_sw_inj_frames`, available under the LALSuite software package [73]. The actual data was stored as SFTs according to the standard SFT_v2

description [70]. Using a duration of 1800s, 9124 and 10124 SFTs were generated from the data collected by the L1 and H1 interferometers, respectively.

The full S6 MDC was performed by five pipelines: *Powerflux*, *SkyHough*, *FrequencyHough*, *Einstein@Home* and *Time Domain \mathcal{F} -Statistic*. Each pipeline established a different tradeoff between robustness and sensitivity, employing different strategies to deal with noise disturbances. Also, different computational capabilities dictate the way in which the parameter space could be treated. The computational implementation of each pipeline can be freely accessed under the LALSuite software package [73].

Figure 4.2 shows a part of the obtained results. For instance, *Einstein@Home*, being a world-wide collaborative project, is able to employ longer coherent times, reaching the deepest sensitivity of all searches; however, its treatment of noise disturbances, which replaces known lines by realizations of Gaussian noise, prevented any detection of signals overlapping with known lines. On the other hand, the *SkyHough* pipeline uses a much shorter coherent time, as already exposed during chapter 3; consequently, it is less sensitive than *Einstein@Home* on the set of known injections. However, its implementation does not *remove* narrow lines; rather, their effects are mitigated by imposing a common threshold over the normalized power. As a result, it achieves a significant efficiency with respect to other pipelines when the search is performed on parameter space regions overlapping with known lines.

4.2.1 The Mock Data Challenge as a Test Dataset

SkyHough results on the Mock Data Challenge serve as a snapshot of the operational state of the pipeline. As so, they can be used to quantify the effectiveness of new data analysis strategies introduced in the pipeline by targeting multiple injections and comparing its results. To do so, we prepared a set of 176 known injections of the challenge, shown in figure 4.3.

The sample of injections is significant in the sense that the whole parameter space of interest (the one which produces frequency modulations) is uniformly covered; that is, frequency presents a uniform distribution, spin-down presents a logarithmic uniform distribution, right ascensions is uniformly distributed and declination is sinusoidally distributed to ensure a uniform coverage of the celestial sphere. The fraction of spin-up signals is consistent with the one used during the MDC. As for the sensitivity depth, a wide range, beyond the one achieved by the *SkyHough* pipeline, is covered.

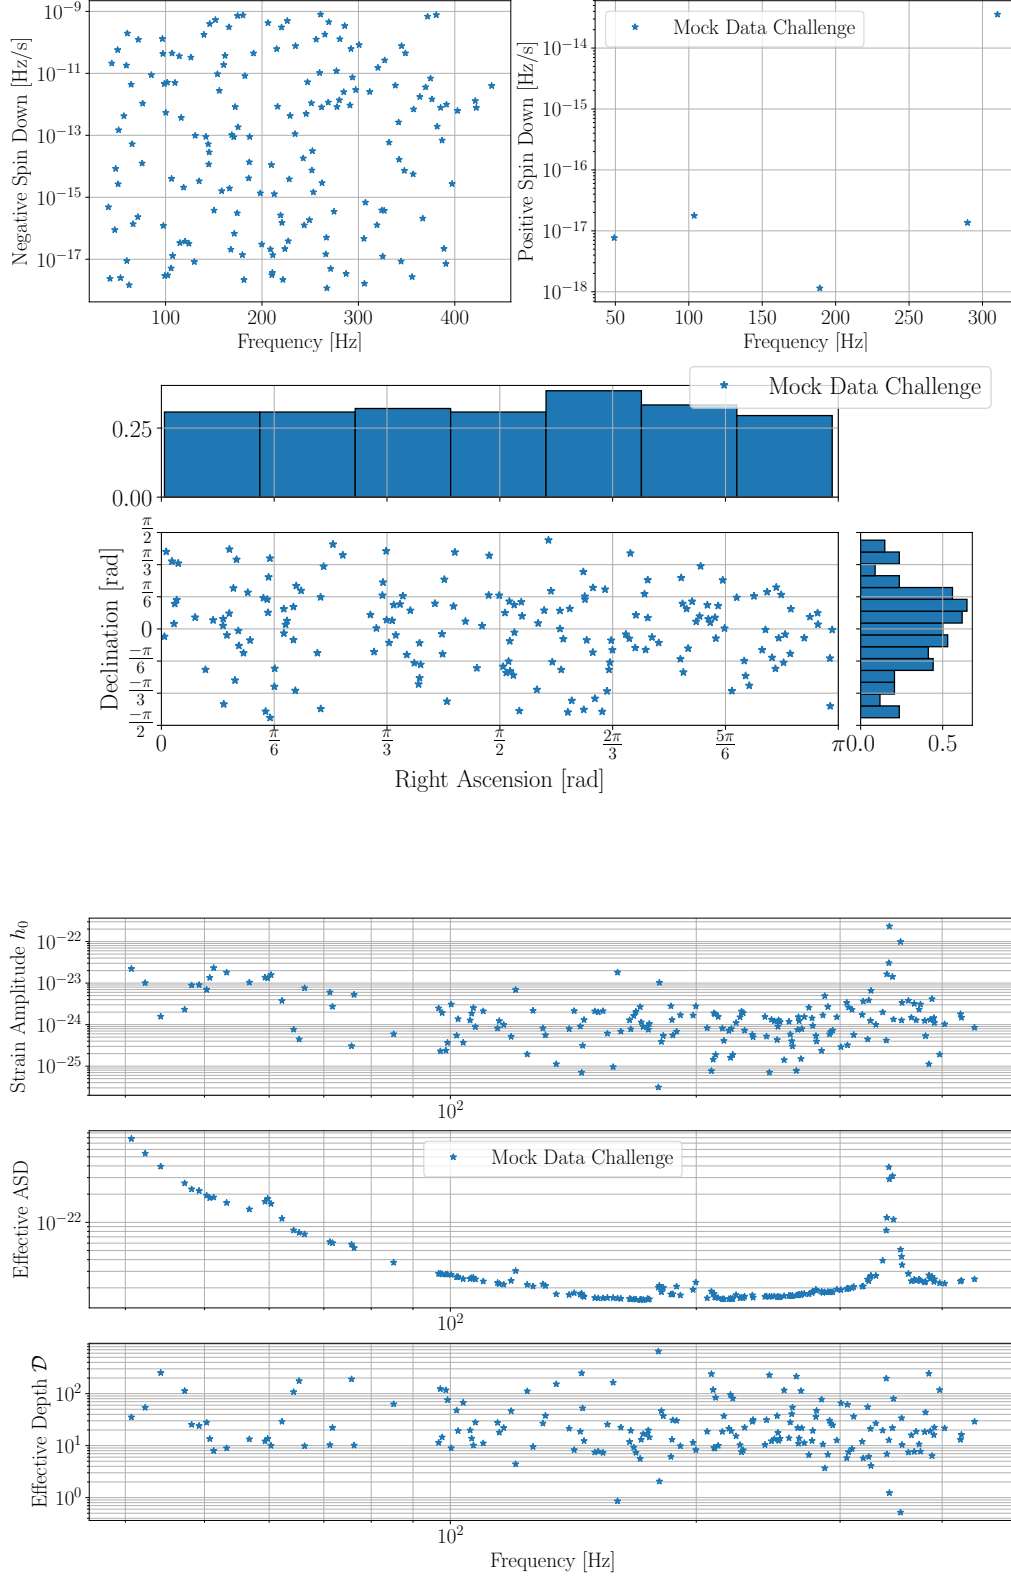


Figure 4.3: Distribution of the employed subset of injections along the parameter space. The injections are uniformly distributed in frequency, spin-down and the celestial sphere. The amount of spin-up injections represents a small fraction of the whole, resembling the intended Mock Data Challenge distribution.

Part II

Original Results

A Noise Robust Statistic for the SkyHough Pipeline

This chapter presents a novel development on the *SkyHough* pipeline, which we label as the *Robust Statistic*. The main idea is to randomly sample the parameter space in order to probe the underlying noise distribution; then, a new *robust* critical ratio can be constructed. That way, candidates' significance will be quoted with respect to the *actual* noise distribution, automatically taking into account variations on the number count noise floor, which nowadays have to be treated *ad hoc* using veto procedures.

Moreover, the suppression of noise contributions into the number count statistic allows for the construction of a combined, multi-interferometer statistic. This idea is applied to the *robust* critical ratio, and increases the sensitivity of the search to signals which were not detected by both detectors using the former, *Gaussian*, critical ratio.

A subset of signals produced for the S6 Mock Data Challenge were used to test the new implementation. Data was stored as 1800s SFT_v2 files [70] in the California Institute of Technology's supercomputing cluster (CIT) [40]. The Hanford interferometer (H1) yielded 10124 SFTs, while the Livingstone interferometer (L1) yielded 9124 SFTs. Its use was motivated by the availability of previous results of the *SkyHough* pipeline, which would ease the evaluation on improvements.

5.1 The Robust Statistic

The former exposition on the *SkyHough* pipeline described the number count statistic for the case of pure noise, yielding as a result a Gaussian distribution. The proof relies in two steps: First, we described the noise number count as a sum of identical independent Bernoulli variables with a certain success probability given by the false alarm probability α ; that is, a Binomial distribution. Then, we invoked the Central Limit Theorem to approximate such a distribution by a Gaussian distribution, the mean and variance of which was directly derived from the previously stated false alarm probability.

The derivation of α in terms of a power threshold assumed a zero-mean, stationary Gaussian noise. That is, noise was modeled as stationary fluctuations over a null-power floor. However, we already discussed during section 4.1 showed the variety of structures which could be generated by noise. In particular, spectral disturbances (such as *lines*) substantially drift away from the proposed assumptions, directly affecting the number count of a wide region of the parameter space. As a result, the *true* number count statistic differs from the Gaussian prescription given by (3.34).

The formal treatment of the Robust statistic goes as follows: Consider the time-frequency track assigned to a particular template $T(\vec{\lambda})$. Instead of an overall, homogeneous probability α , each bin $(a, k) \in T(\vec{\lambda})$ has a certain selection probability α_{ak} due to the non-homogeneity of the noise. That is, different frequency components are affected by different types of noise at different times. Therefore, the number count due to noise fluctuations can be approximated by a sum of non-identical, independent Bernoulli random variables.

Such a description results in a Poisson-Binomial distributed variable. The amount of Bernoulli trials involved in the number count computation is given by the track length, which is related to the number of SFT in used. Depending on the observing run, $\mathcal{O}(10^{3-4})$ are involved in the analysis. As for the success probability of the trials, it is given by the false alarm probability, which is never close to 1 or 0. Hence, from an analytical point of view, it can be shown that the Poisson-Binomial distribution is well represented by a Poisson distribution with the suitable event rate parameter [49]. Moreover, the typical values of the Poisson's event rate parameter, which are related to the average number count noise value, lead the Poisson distribution to the regime in which the Central Limit Theorem becomes useful; hence, the Gaussian assumption on the number count noise distribution is still valid under this new formulation of its statistics, yet the mean value and variance differ from the ones quoted in (3.34).

This argument is intended to justify the validity of the statistical arguments developed in [46]. Nevertheless, such conditions do not have to be met during a real data analysis; hence, the underlying noise statistic will not be derived from this procedure; rather, we will probe the number count noise distribution by means of a parameter space sampling.

Figure 5.1 exemplifies the idea behind the parameter space sampling. Two number count distributions were generated using a Gaussian distribution. The first one corresponded to a noise number count, while the second one corresponded to a noise distribution under the presence of an astrophysical signal.

As explained during chapter 3, astrophysical signals *rise* the number count distribution according to a non-centrality parameter λ . If we focus on an ideal situation, shown in upper figure 5.1, noise number count fluctuates around the null value, while the presence of a signal increases the mean value of the number count distribution according to the non-centrality parameter $\lambda_s \neq 0$.

However, this situation can be easily broken by spectral disturbances. For instance, we could be working under the presence of a line. Lines, as sustained disturbances, can be described in terms of an associated non-centrality parameter $\lambda_n \neq 0$ which rises the average number count fluctuation, as shown in lower figure 5.1. The use of a Gaussian number count statistic, which assumes $\lambda_n = 0$, yields disastrous results, as the significance of a pure noise sample is actually *overestimated*.

On the other hand, the estimation of noise statistics using sampled values probes the actual statistic of that particular region of the parameter space. As a result, noise and signal significance are properly stated. It is worth mentioning that this procedure is compatible with the previous statistical prescription, as the sample statistic coincides in an ideal case. This is shown in upper 5.1 as well.

Summarizing, we will use a sampling procedure to probe the actual noise number count statistic. The result will be a sample mean μ_r and a sample standard deviation σ_r for a certain region of the parameter space. Then, templates within that region of the parameter space will get assigned a *robust* critical ratio Ψ_r according to its number count n in a consistent manner with the previous definition of critical ratio

$$\Psi_r = \frac{n - \mu_r}{\sigma_r}. \quad (5.1)$$



Figure 5.1: Histogram of the Gaussian and robust statistics for a simulated set of number counts under different noise conditions. Blue histograms represent the number count distribution produced by background noise; yellow histograms represent the number count distribution in the presence of a signal. Noise’s characteristics are stated in terms of the non-centrality parameter (3.25), as explained in the main text. **Upper:** Number count simulation using zero-mean Gaussian noise. **Lower:** Number count simulation in the presence of a persistent spectral disturbance, modeled by a non-zero λ_n .

5.2 The Multi-Interferometer Statistic

Gravitational wave searches are performed with multiple interferometers (IFOs) at once. Hence, it is natural to consider ways of combining single-interferometer statistics into a multi-interferometer statistic. Developments in this matter can be traced back to [41] for the case of coherent analysis. We will extend those ideas towards a multi-interferometer critical ratio, the antecedents of which can be found in [59].

The inclusion of multiple interferometers into the analysis potentially increases the significance of gravitational wave candidates, given that statistical significance is jointly added; however, spurious noise fluctuations could also become more significant, increasing the false alarm probability of the search. Through the use of the robust statistic, we expect to suppress noise artifacts present in our data, effectively countering the increase of false alarm in our multi-interferometer statistic.

We will combine the robust critical ratio from different interferometers; that is, each interferometer will compute its own robust statistic, delivering a robust critical ratio for each template. Then, critical ratios will be mixed using the suitable measure data quality measurement of each detector.

An initial choice was made in [59], where weights (3.8) were taken to measure the data quality of a detector as they explicitly take into account the main data degradation processes: Noise floor and detector response. Let us define $\text{IFO} = \{I_1, \dots, I_M\}$ as the set of interferometers to work with. Each interferometer $I \in \text{IFO}$ holds a set of weights \vec{w}_I according to chapter 3. Then, the exact data quality measurement for the I -th interferometer was taken to be

$$r_I = \frac{\sqrt{\|\vec{w}_I\|^2}}{\sqrt{\sum_{J \in \text{IFO}} \|\vec{w}_J\|^2}}, \quad (5.2)$$

which can be re-expressed in terms of the single-interferometer number count standard deviation as

$$r_I = \frac{\sigma_I}{\sqrt{\sum_{J \in \text{IFO}} \sigma_J^2}}. \quad (5.3)$$

Thus, we will describe the multi-interferometer critical ratio as follows: For a particular template, each interferometer $I \in \text{IFO}$ will deliver a number count n_I , which will get associated a mean value μ_I and a standard deviation σ_I using the parameter space sampling. Assuming null inter-interferometer correlation, the multi-interferometer number count n_m and its corresponding statistics will be defined as

$$n_m = \sum_{I \in \text{IFO}} n_I; \quad (5.4a)$$

$$\mu_m = \sum_{I \in \text{IFO}} \mu_I; \quad (5.4b)$$

$$\sigma_m^2 = \sum_{I \in \text{IFO}} \sigma_I^2. \quad (5.4c)$$

As a result, the multi-interferometer critical ratio is given by

$$\begin{aligned} \Psi_m &= \frac{n_m - \mu_m}{\sigma_m} \\ &= \sum_{I \in \text{IFO}} \Psi_{rI} \frac{\sigma_I}{\sigma_m}, \end{aligned} \quad (5.5)$$

where Ψ_{rI} represents the robust critical ratio associated to the I -th interferometer and the multi-interferometer weights are consistent with 5.3.

[59] obtained information on the detector state using a prescribed quantity, namely the weights; we supply such information in terms of a parameter space sampling, taking into account the *actual* noise distribution.

5.3 Implementation of the Robust Statistic

Both procedures were added to the `DriveHoughMulti.c` code, which encodes the computational core of the *SkyHough* pipeline. We proceed to explain the way in which they were implemented. It is important to recall that previous functionalities are by no means overridden by new implementations; that is, we are still capable of performing data analysis according to former prescriptions.

Codes are publicly available under the LIGO Algorithm Library LALSuite [73]. The implementation made use of the GNU Scientific Library (GSL) [30]. *Valgrind* [58] was used as a debugging tool during the development and test processes.

5.3.1 Parameter Space Sampling

The *SkyHough* pipeline segments the parameter space in a particular fashion: Chosen a frequency band, gravitational wave signals are looked for in different sky locations (*sky patches*) for a certain range of spin-down parameters. Hence, we decided to perform a parameter space sampling for each sky patch, for each frequency band, in order to compute the corresponding statistic. This procedure is done for each interferometer, allowing us to compute the multi-interferometer statistic at the same time.

The implementation of the sampling procedure benefits from a previous code, intended to calibrate the number count χ^2 statistic [24]. From an operational point of view, independent samples are drawn for each of the four parameters within the established parameter space limits. That is, a random frequency is selected within the considered frequency band; a random sky position is selected within the studied sky patch; a random spin-down value is selected from the given spin-down range. Then, the number count associated to the randomly selected template is computed, taking the suitable weights into account.

The effects of a continuous wave signal into the parameter space sampling can be argued as follows: Astrophysical signals, opposed to wide spectral disturbances, are expected to affect a much more confined region of the parameter space; hence, they are less likely to affect the noise estimation performed by a random number count sampling.

5.3.2 Estimation of Sample Statistics

The result of a sample was used to estimate the mean value and standard deviation of the underlying number count noise distribution. However, the usual computational recipes to estimate the standard deviation of a sample can be easily affected by *catastrophic cancellations*, since they involve the subtraction of floating point numbers way bigger than the final value. Moreover, if we intended to implement another strategy, such as re-sampling method to ensure a good representation of the parameter space, we would need to deal with a significant amount of numbers held in memory at once.

Both issues can be overcome with the use of Welford's Algorithm [81]. It is described as an *online* algorithm, which means that the estimation is performed as samples are taken; i.e. there is no need to keep every single sample in memory to perform the statistical estimation. Moreover, it allows the variance computation to be written in a way such that catastrophic cancellation is avoided.

Given a set of n samples $\{\xi_1, \dots, \xi_n\}$, the algorithm describes the estimation of the mean μ and standard deviation σ as a recursive relation for $m = 1, \dots, n$:

$$\mu_m = \mu_{m-1} + \frac{\xi_m - \mu_{m-1}}{m}, \quad (5.6)$$

$$M_m = M_{m-1} + (\xi_m - \mu_{m-1})(\xi_m - \mu_m), \quad (5.7)$$

where the sample mean is given by $\mu = \mu_n$ and the sample variance is given by $\sigma = M_n/(n-1)$. As clearly seen from both expressions, the subtractions are always performed among numbers of the same order, reducing the precision loss with respect to a brute force approach.

The combination of this algorithm with the previous sampling procedure yields an effective code which practically does not increase the computational load with respect to the previous version, as there is no need to spend significant amounts of memory to store new parameters.

5.3.3 Multi-Interferometer Number Count

The final modification with respect to the previous code regards the computation of number counts. The new prescription for the multi-interferometer number count statistic 5.5 requires to individually compute every single-interferometer critical ratio. To do so, the internal data management of the code had to be modified, as the interferometer information from a multi-interferometer data set was lost due to the procedure implemented in the previous code.

After our modification, every set of SFTs coming from a different interferometer is stored as a block with a certain order. The initial and final SFT's indexes of each block is kept in memory in order to track every single-interferometer data set during the number count computation process. As a result, we are able to

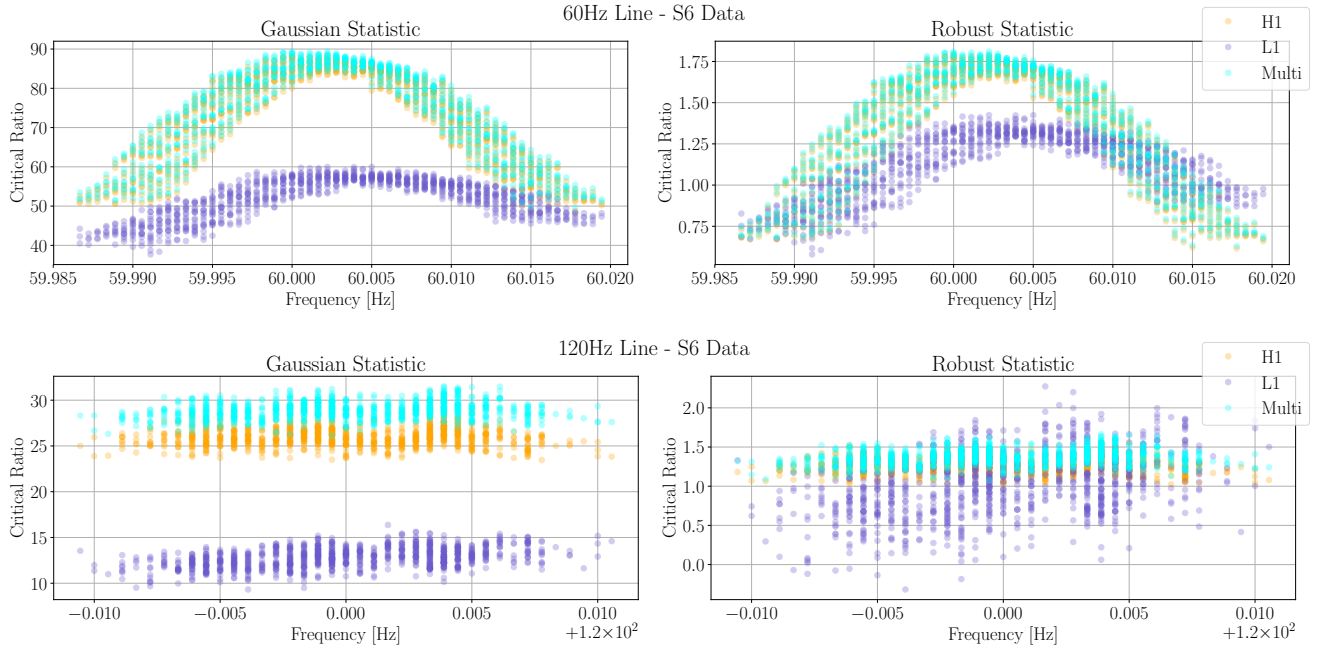


Figure 5.2: Critical ratio comparison on the 60Hz and 120Hz harmonic of the detector’s AC power supply during the S6 science run. The analysis kept the best 1000 candidates according to its multi-interferometer statistic. **Left:** Single and multi-interferometer statistics constructed using the Gaussian *SkyHough* Critical Ratio. **Right:** Single and multi-interferometer statistics constructed using the Robust Critical Ratio.

separately compute single-interferometer number counts and combine them according to 5.5, which can be considered as an enhancement with respect to the previous version of the code.

5.4 Performance of the Robust Statistic

We used the novel robust, multi-interferometer statistic over two situations in order to study its behavior and capabilities regarding the suppression of noise artifacts and the detection of continuous gravitational waves. The former item was tested by pointing the algorithm towards parameter space regions populated by known spectral disturbances; the later used a subset of injections from the S6 Mock Data Challenge to compare its sensitivity with respect to the previous implementation of the *SkyHough* pipeline.

Computational jobs were executed in the supercomputer held at the California Institute of Technology (CIT) [40], using HTCondor [69] as the job manager.

5.4.1 Analysis of Known Lines

We performed a search for continuous gravitational wave searches on regions of the parameter space that were known to be affected by spectral disturbances. We pointed the sky location towards the polar region, which is where Doppler shift is less severe. This implies that lines yield similar patterns to those of a signal, as they do not suffer a frequency modulation either.

Figure 5.2 shows the critical ratio obtained after searching over the frequency band poisoned by the 60Hz and 120Hz lines. These lines arise from the very functional core of the interferometer, since they are

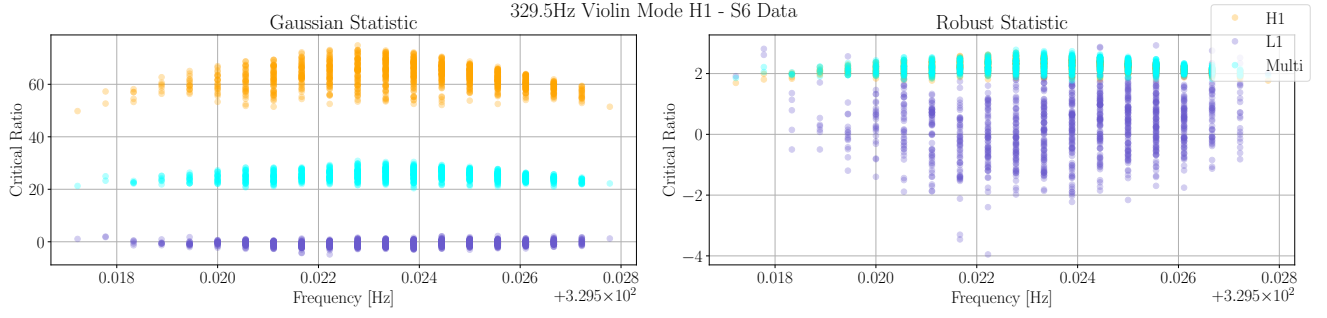


Figure 5.3: Critical ratio comparison for a violin mode present in the H1 interferometer during the S6 science run. The analysis kept the best 1000 candidates according to its multi-interferometer statistic. **Left:** Single and multi-interferometer statistics constructed using the Gaussian *SkyHough* Critical Ratio. **Right:** Single and multi-interferometer statistics constructed using the Robust Critical Ratio.




Figure 5.4: Critical ratio comparison for a narrow line feature present in the L1 interferometer during the S6 science run. The analysis kept the best 1000 candidates according to its multi-interferometer statistic. **Left:** Single and multi-interferometer statistics constructed using the Gaussian *SkyHough* Critical Ratio. **Right:** Single and multi-interferometer statistics constructed using the Robust Critical Ratio.

related to the AC power supply frequency used in the U.S. As we can see, the robust statistic suppresses the influence of such artifacts to a lower value; this is precisely the expected behavior, as the *virtual* significance of the candidates given by the Gaussian statistics is due to the increased value of the mean noise number count.

Figure 5.3 shows an equivalent result for a different line. In this case, we took a violin mode which was specially prominent in the LIGO Hanford interferometer (H1). Such modes arise from the vibrations of mechanical suspensions of the interferometer mirrors. They pose a bigger challenge than the 60Hz harmonics, as their monitoring is inherently more difficult. Nevertheless, the probing strategy of the robust statistic is able to suppress its influence for the very same reason as before.

So far, the highlighted spectral disturbances span across relatively wide regions of the parameter space; this is the desirable setup for the application of the robust statistic and, as already shown, performs in a competent way against the previously prescribed Gaussian statistic.

For the sake of completeness, we performed a last search over a narrow line feature on the LIGO Livingston detector (L1). The performance of the robust statistic over this particularly narrow line delivers us valuable information on the effects of the robust statistic over gravitational wave signals: As previously commented, gravitational wave signals affect a relatively small region of the parameter space, and produce indistinguishable traces compare to that of persistent noise disturbances when a search is performed close to the polar regions.

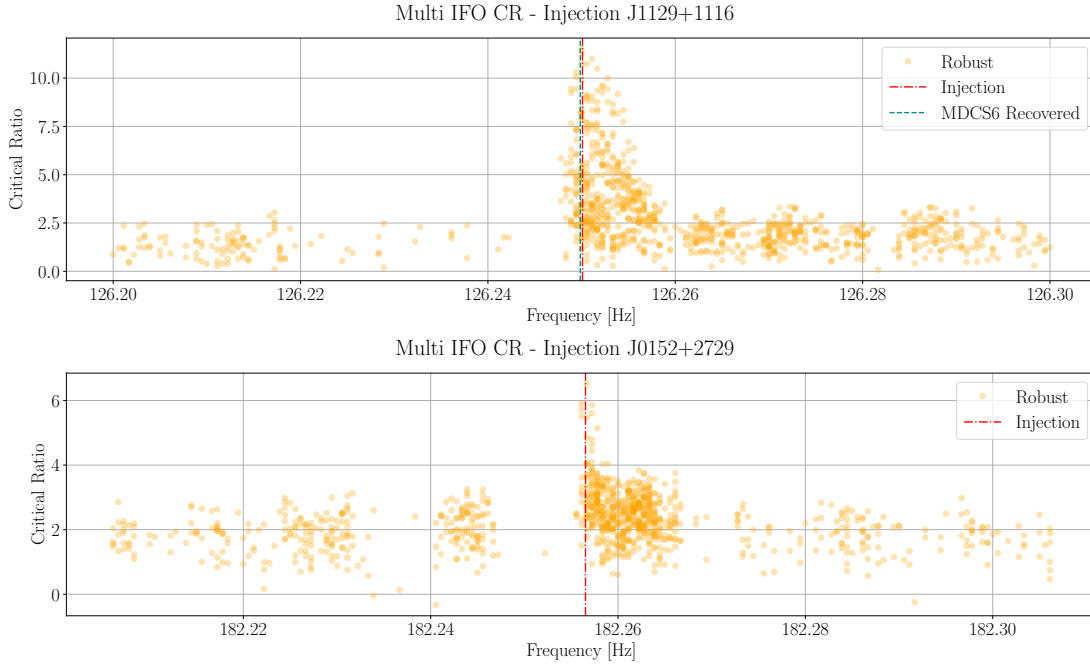


Figure 5.5: Results of the continuous wave search on two injections of the Mock Data Challenge S6. The 1000 most significant candidates were kept, according to its multi-interferometer critical ratio. The vertical red line represents the injection point; the vertical purple line represents the recovered parameter by the SkyHough pipeline during the Mock Data Challenge. **Upper:** Injection recovered by the *SkyHough* pipeline and the multi-interferometer robust statistic. **Lower:** Injection *not* recovered by the *SkyHough* pipeline during the Mock Data Challenge S6, but recovered by the multi-interferometer robust statistic according to the preliminary post processing.

The results are shown in figure 5.4. Significance mitigation is present, as the noise number count floor has been raised due to the use of a parameter space sampling; however, the suppression is significantly lower than in the previous cases. This fact does *not* discard the robust statistic for continuous gravitational wave searches, as the significance of a small regions of the parameter space stands still. In fact, one could attempt to construct a veto procedure to discriminate between astrophysical signals and spurious noise fluctuations according to its critical ratio mitigation.

Summing up, we tested the basic feature of the robust statistic: Noise contributions are suppressed thanks to the use of accurate background statistics, obtained by means of parameter space samplings; regarding narrow perturbations, its mitigation is significantly lower, allowing us to use this idea to search for continuous gravitational waves.

5.4.2 S6 Mock Data Challenge: Overview

After understanding the main traits of the new robust multi-interferometer statistic, we used the already introduced Mock Data Challenge to directly compare the results with the previous formulation of the *SkyHough* pipeline. The subset of the data set, as well as its fairness in terms of signal distribution in the parameter space, were already introduced during chapter 4.

Regarding parameter space resolution, coherent time was $T_{\text{coh}} = 1800\text{s}$, while total observation time was $T_{\text{obs}} \simeq 4 \cdot 10^7\text{s}$, as quoted during chapter 4. Hence, the resulting parameter space resolutions where

$$\delta f \simeq 5 \cdot 10^{-4}\text{Hz},$$

$$\delta f_1 \simeq 1 \cdot 10^{-11} \text{Hz/s}.$$

As for the sky resolution, which depends on the frequency according to (3.18), a pixel factor $p = 2$ was taken.

The searches were performed around the injection points of the parameter space. That is, the analysis focused on a 0.1Hz frequency band centered around the injection point, pointing towards the sky position of the injection. For each injection, we selected the 1000 most significant candidates according to the multi-interferometer robust statistic. Then, we imposed a critical ratio threshold $\Psi_{\text{th}} = 5$, disregarding any candidate with a lower critical ratio. The final estimation of signal parameters was done using a weighted average among the surviving candidates within each frequency band, using its multi-interferometer robust critical ratio as weights. This post processing step is meant to obtain a preliminary result on the sensitivity of this new statistical approach; future studies on the subject of multi-interferometer shall improve it in terms of new prescriptions and veto procedures in order to diminish the false alarm of a search, as current post processing steps based in coincidences do.

The fairness of the preliminary post processing step is shown in figure 5.5, where two particular injections are shown. The first one was recovered by the both procedure, the *SkyHough* pipeline and the multi-ifo robust statistic using the preliminary post processing step. As for the second one, only the new statistic is able to recover it.

This detection can be related to the combined action of both interferometers into the search, as opposed to the strategy employed during the Mock Data Challenge, which combined the information of each detector *after* the execution of the Hough transform. This, in turn, is made possible by the use of the robust statistic, which mitigates the combined noise distribution introduced by both detectors, as clearly seen in both figures 5.5, where the surrounding background critical ratio is consistently kept below the 4σ value.

As for the results, an outlook is shown in figure 5.6, which includes the injections detected by both implementations of the pipeline, namely the one used during the Mock Data Challenge and the robust multi-interferometer statistic with a preliminary post-processing. 127 out of 176 injections were successfully recovered; 11 of those injections were not recovered by the previous formulation of the *SkyHough* pipeline.

Every injection detected during the Mock Data Challenge S6 was recovered by the new implementation as well. This fact is consistent with our arguments, given that the robust statistic gauges the noise contribution: If an injection was already significant with respect to the background *before* the implementation of the robust statistic, It will be still significant after its implementation.

Newly detected injections are well distributed along the parameter space, with no particular preference for a particular sky position or frequency band. This ensures that the parameter space sampling is representative. For instance, if the recovered injections were located at the polar regions, it would imply that the robust statistic is not suitable to treat with the effects of the Doppler shift. However, the number count computation already takes such effect into account; hence, the use of parameter space samplings abstracts the Doppler shift out of the robust statistic.

Looking into the depth distribution, we clearly see how the robust statistic is able to recover injection *further* into the depth distribution. That is, there is no apparent depth gap between the originally recovered injections and the newly recovered ones. Following the previous argument, this is the expected behavior, since the robust statistic treats every region on the parameter space in the same manner.

In brief, the results delivered by the robust statistic are consistent. Every injection recovered by the previous formulation is recovered by the robust statistic; this fact can be related with the argument given to figure 5.1: The robust statistic refers the significance to the *actual* noise distribution; hence, those

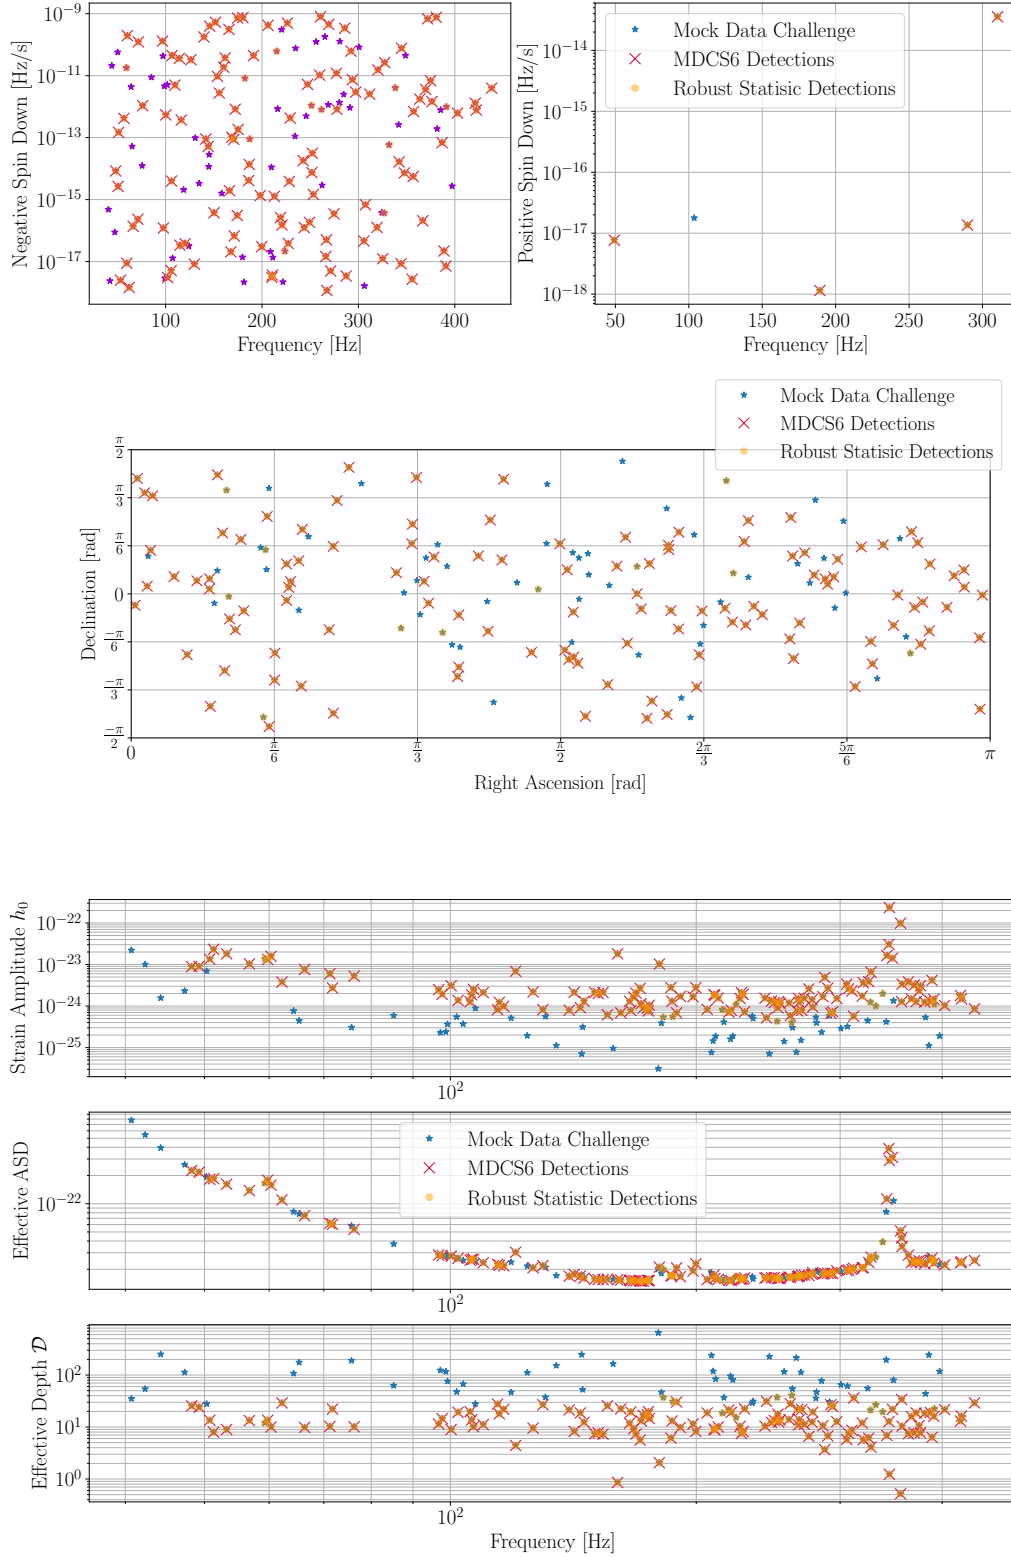


Figure 5.6: Direct comparison of Mock Data Challenge signal detections between both formulations of the *SkyHough* pipeline. Stars represent injections of the S6 Mock Data Challenge. Crosses mark those injections which were recovered by the previous formulation of the *SkyHough* pipeline. Circles mark injections recovered by the multi-interferometer robust statistic according to the preliminary post processing. This plot is equivalent to figure 4.3, found in chapter 4.




Figure 5.7: Efficiency curve for the Mock Data Challenge results of the SkyHough pipeline and the novel multi-interferometer robust strategy with the preliminary post-processing. The vertical axis represents binary values for (non-)detected injections and the frequentist detection probability for the sigmoid fit. The horizontal axis represents the log inverse depth $\log_{10} \mathcal{D}^{-1}$.

injections which are prominent enough over the noise stand still after the use of the robust statistic. As for the newly recovered injections, they are due to the use of the multi-interferometer statistic. As previously discussed, we are able to use this particular strategy thanks to the characterization of the background noise and ensuing mitigation of noise artifacts performed by the robust statistic.

5.4.3 Efficiency and Parameter Recovery

We proceed to compare the efficiency of both approaches using a standard metric within the continuous wave searches: The 95% detection probability depth [8, 9]. It corresponds to the sensitivity depth at which 95% of signals are expected to be recovered in a frequentist sense. The use of that magnitude instead of a gravitational wave amplitude h_0 abstracts the dependency on the surrounding amplitude spectral density, easing the computation of other derived quantities, such as upper limits.

Usually, several signals are injected at different sensitivity depths in order to estimate the detection probability. Then, a sigmoid function is used to interpolate the obtained probabilities for a certain range of depths. If the procedure is properly done, the depth corresponding to a 95% detection probability $\mathcal{D}^{95\%}$ will be obtained as a result; then, one can propagate other physical properties converting the sensitivity depth to an upper limit over h_0 using (1.12). It is worth to note that the sigmoid fit is performed over the sensitivity depth (or the characteristic amplitude h_0) using a logarithmic scale [8, 9].

However, the idiosyncrasy of our analysis prevents the existence of a significant number of injections at the same sensitivity depth, as they were produced with a different purpose in mind [80]. Hence, we will use the sigmoid function approach in terms of a decision boundary, interpolating such boundary after the labeling of injections as detected/not detected.

To do so, we compute the sensitivity depth \mathcal{D} of each injection using (1.12) with the effective noise amplitude spectral density, as defined in chapter 4. Then, each injection is labeled as *detected* (1) if it survives to the preliminary post processing step, being otherwise labeled as *not detected* (0). Finally, a sigmoid function is fitted to the data, using the logarithm of the inverse depth for that purpose. For a desired detection probability c , the fit is inverted to yield the corresponding sensitivity depth \mathcal{D}^c .




Figure 5.8: Parameter space projected distance (3.45) between the injection points and the recovered parameters by both formulations of the *SkyHough* pipeline. The vertical axis represents the total parameter space distance, taking into account the combined distances within the three subspaces and using bins as units. The horizontal axis represents the log inverse depth $\log_{10} \mathcal{D}^{-1}$.

The sigmoid function is defined as

$$f(x; a, b) = \frac{1}{1 + e^{-(ax+b)}}. \quad (5.8)$$

The fit is performed using SciPy library's `curve_fit` function [45]. As for the statistical uncertainty, we compute the 1σ envelope of the fit in terms of its covariance matrix C as

$$\sigma_f = \pm \sqrt{\left(\frac{\partial f}{\partial a}\right)^2 C_{aa} + \left(\frac{\partial f}{\partial b}\right)^2 C_{bb} + 2 \left(\frac{\partial f}{\partial a}\right) \left(\frac{\partial f}{\partial b}\right) C_{ab}}, \quad (5.9)$$

where the subindex on C represent the elements of the covariance matrix corresponding to each pair of parameters and the dependencies of f are kept implicit.

We used the logarithm of the inverse depth $\log_{10} \mathcal{D}^{-1}$ to perform the fit, which is proportional to the logarithm of the gravitational wave amplitude $\log_{10} h_0$; that is, the more negative the depth inverse logarithm, the fainter the signal with respect to the surrounding noise.

The result of this procedure, which is done for both the Mock Data Challenge results and the new statistic *in the same manner*, is shown in upper figure 5.7. The fitted parameters were

$$\begin{cases} a_{\text{MDCS6}} = 6 \pm 1 \\ b_{\text{MDCS6}} = 20 \pm 4 \end{cases} \quad (5.10)$$

for the formulation of the *SkyHough* pipeline used during the S6 Mock Data Challenge, while the new formulation obtained

$$\begin{cases} a_r = 6 \pm 1 \\ b_r = 22 \pm 3 \end{cases}. \quad (5.11)$$

This results are propagated to the estimation of the 95% detection probability depth $\mathcal{D}^{95\%}$, obtaining

$$\mathcal{D}_{\text{MDCS6}}^{90\%} = 18.2 \pm 0.5 \quad (5.12)$$

for the initial implementation of the *SkyHough* pipeline and

$$\mathcal{D}_r^{90\%} = 21.6 \pm 0.4 \quad (5.13)$$

for the new statistical formulation.

We obtain an increase $\sim 20\%$ of the sensitivity depth by using the novel multi-interferometer robust statistic with a preliminary post processing. This result acts in favor of this new approach, as it represents a significant increase in sensitivity at a very low computational cost, as argued during previous section.

But, we need to deal with a last issue in order to clarify this new development as a valid approach to detect continuous waves: The parameters related to the detected injections need to be *close* enough to be useful for the forthcoming follow-up steps. To quantify the accuracy of the recovered parameters, we compute its parameter space distance to the injection point using the projected prescription (3.45), introduced during chapter 3.

Figure 5.8 shows the distance to the injection point delivered by both prescriptions of the *SkyHough* pipeline. We represent the total distance in terms of parameter space bins. Common injections are recovered at a similar deviation, in a consistent manner with the previous discussion. Newly recovered injections present a consistent deviation from the injection point with respect to previously detected injections.

This result confirms the effectiveness of the robust multi-interferometer statistic, as every injection recovered by the initial formulation of the *SkyHough* pipeline is recovered with a similar deviation by the new statistic and the deviation yielded by the injections solely detected by the new statistic are consistent with the common bulk of signals.

Summarizing, the novel statistical formulation of the *SkyHough* pipeline has proven to be a feasible improvement towards the detection of continuous gravitational waves. The use of a parameter space sampling allows us to combine multiple detectors in a consistent manner, lowering the appearance of noise artifacts thanks to the properly gauge background contributions.

We recall that this results are preliminary, as they conform a subset of injections over which a simplified post-processing has been applied. An extended analysis to legitimate this approach could use state-of-the-art data to study the appearance and distribution of outliers or a re-computation of upper limits to gauge the actual improvement on sensitivity.

Future Work

During the development of the present work, several tools were tempted as interesting strategies to be ported to the *SkyHough* pipeline. Although they are still in a very early stage of experimentation, it is worth mentioning them, as they influenced some of the basic ideas described in this work.

This chapter will expose three different data analysis techniques that target different stages of the detection pipeline. Each of them could represent a significant improvement in the way data is analyzed during a search for continuous gravitational wave searches. The first section will expose the *tau statistic*, an alternative formulation of the coherent step of a semicoherent search; section two will address the issue of upper limits, using the *universal statistic* to estimate them using the statistical properties of a sample of data; section three will briefly comment a reformulation of the post processing step in terms of a machine learning algorithm to ease the calibration process.

6.1 Tau Statistic

The tau (τ) statistic [19] corresponds to a generalized formulation of the coherent step presented during section 3.1.1. We proceed to review such step in order to point out its problems and the way the τ -statistic addresses them.

Consider a set of N samples from a sinusoidal signal with frequency f , taken at different times $t_j \in [0, T_{\text{coh}}]$, where $t_j = t_0 + j\Delta t$ for a suitable Δt in order to fulfill the Shannon-Nyquist sampling condition. A sample taken at time t_j will be labeled as $x[j]$; thus,

$$x[j] \sim \cos(2\pi f t_j + \phi), \quad (6.1)$$

where irrelevant amplitude constants haven been omitted.

Following the coherent step, we compute the discrete Fourier transform of our samples, given by $\tilde{x}[k]$. As stated during chapter 3, $[k]$ refers to the k -th frequency component, which is discretized in units of T_{coh}^{-1} .

So far, nothing differs from the previously specified coherent step. However, let us explicitly write the discrete Fourier transform of the sampled sinusoidal. Expressing the signals' frequency in terms of the discretization $f = lT_{\text{coh}}^{-1}$, we obtain

$$\tilde{x}[k] \sim e^{i\phi} D_N(k-l) + e^{-i\phi} D_N(k+l), \quad (6.2)$$

where D_N represents the *Dirichlet* kernel, given by

$$D_N(z) = e^{i\pi z(1-\frac{1}{N})} \frac{\sin(\pi z)}{N \sin(\frac{\pi z}{N})}. \quad (6.3)$$

This function is peaked at $z = 0$. For $z \in \mathbb{Z}$, it can be proven that $D_N(z) = \delta_{0z}$, where δ_{ij} represents the Kronecker delta, which returns one for $i = j$ and zero for $i \neq j$; otherwise, the Dirichlet kernel falls off as a power law $\sim z^{-1}$.

For this reason, we can impose a simplification into (6.2) given by $D_N(k+l) \ll D_N(k-l)$. Moreover, the usual number of samples involved in a search is big enough to consider $N^{-1} \ll 1$. Finally, the rapid fall off allows us to only consider nearby frequency bins; this conditions be stated in terms of a Taylor expansion of the sinusoidal function $N \sin(\frac{\pi z}{N}) \sim \pi z$. Hence, the Fourier transform of a sampled sinusoidal with intrinsic frequency $f = lT_{\text{coh}}^{-1}$ can be expressed as

$$\tilde{x}[k] \sim e^{i(\phi+\pi(k-l))} \text{sinc}(k-l), \quad (6.4)$$

where sinc states for the cardinal sine function. Here resides the critical step: If we consider $l \in \mathbb{N}$, that is, f is one of the *resolved* frequencies of the discrete Fourier transform, the signal is entirely contained within a frequency bin and, as a result, $|\tilde{x}[k]|^2$ represents the optimal statistic to detect the presence of a signal, according to the Neyman-Pearson criterion. This approach is the one taken by the *SkyHough* pipeline, as shown during chapter 3.

However, if f is *not* one of the resolved frequencies, i.e. $f \notin \mathbb{N}$ (we consider $f > 0$ without loss of generality), the sinusoidal signal gets *spread* across several frequency bins¹. Hence, the optimal statistic to detect the presence of a signal is not simply the power density of a certain component; rather, one has to *reconstruct* the spread power to obtain such a statistic.

Even though the signal gets spread across the whole frequency spectrum, it can be proven that the neighborhood of a particular bin is enough to recover a significant result. Let us consider a particular frequency bin k and its P nearest neighbors at each side; that is,

$$\tilde{x}[k]_P = \begin{pmatrix} \tilde{x}[k-P] \\ \vdots \\ \tilde{x}[k] \\ \vdots \\ \tilde{x}[k+P] \end{pmatrix}. \quad (6.5)$$

Then, the optimal statistic, which is reconstructed from the data, can be defined for a particular frequency component k as the following bilinear form

$$\tau[k] = \tilde{x}[k]_P^\dagger M_P \tilde{x}[k]_P, \quad (6.6)$$

where M_P represents a symmetric $P \times P$ matrix, defined in terms of the zeroth order spherical Bessel function j_0 as

$$M_P r s = (-1)^{r-s} \int_{-1/2}^{1/2} d\delta j_0[\pi(r+\delta)] j_0[\pi(s+\delta)], \quad (6.7)$$

where $r, s = -P, \dots, 0, \dots, P$ label the matrix components.

The use of this new statistic implies the establishment of the corresponding threshold τ_{th} either in terms of the Neyman-Pearson criterion or maximizing the critical ratio, as done during chapter 3. Note that $\tau_{\text{th}} = \tau_{\text{th}}(P)$, as the reconstructed power depends explicitly on the amount of neighbor frequency bins taken into account. Following an eigenvalue analysis on M_P , it is clear that taking P further than 3 or 4 frequency bins adds irrelevant contributions to the τ -statistic.

¹This phenomenon is sometimes referred to as *spectral leakage*.

6.1.1 Implementation

We implemented a computation for the τ -statistic as an enhancement of the `ConvertToSFTv2.c` code of the LIGO Algorithm Library [73]. The initial version of the code is used to extract particular frequency bands from SFT files in order to reduce the amount of irrelevant information during certain kinds of searches. After our modification, the τ -statistic can be automatically computed on any set of SFTs as they are being extracted.

The code is being tested and debugged in order to ensure the proper implementation of the statistical procedure.

6.2 Universal Statistic

The universal statistic procedure [26] is an algorithm to state conservative frequentist upper limits on a given set of samples. It has been used by the *Powerflux* pipeline to establish upper limits on the all sky search of the S6 and O1 observing runs [77, 8].

The algorithm uses statistical information from data in itself to construct the suitable quantile of the underlying noise distribution. Once this quantile is constructed, the upper limit can be quoted in terms of the most significant outlier.

Given a sample of measurements $\{\xi_1, \dots, \xi_n\}$, the $1 - \epsilon$ confidence upper limit can be constructed as

$$\text{UL}^{1-\epsilon} = \max_i \xi_i - \xi_\epsilon, \quad (6.8)$$

where ξ_ϵ stands for the ϵ -quantile of the distribution underlying our samples. However, the estimation of ξ_ϵ turns out to be computationally difficult, as it requires a sorting algorithm. Moreover, the actual estimation could be biased, leading to a non-significant upper limit.

To solve both problems, a new conservative quantile is constructed using the Markov inequality

$$\mathcal{P}(|\xi| \geq a) \leq \frac{\langle |\xi| \rangle}{a}, \quad (6.9)$$

where ξ is a random variable, $a \in \mathbb{R}$ and \mathcal{P} represents the probability function. This inequality holds for *any* positive random variable; hence, it is completely licit to express it in terms of a positive function $|\xi| \rightarrow f(\xi)$. This kind of transformation allows us to engineer an *effective* quantile using the mean value and standard deviation of the available data. The actual construction uses the Markov inequality to probe the *gaussianity* of the distribution's tails: If the tail of the sample distribution can be overwhelmed by that of a Gaussian distribution, the use of a Gaussian quantile suffices our purposes; if, on the contrary, the amount of samples located within the tail region surpasses the expected amount for a Gaussian distribution, we would use f in conjunction with the constraint imposed by 6.9 to *push* the Gaussian quantile further until it encloses the required amount of data.

The contribution of this algorithm lies in the use of data samples to construct an upper limit, rather than requiring the actual estimation of the detection rate in terms of a set of Monte Carlo injections. Nevertheless *Powerflux* possesses a particular advantage to apply this method: As opposed to the Hough Method, the *Powerflux* statistic can be analytically related to the amplitude of a gravitational wave. That is, once an upper limit has been set to the power samples, it can be automatically transported to the gravitational wave amplitude.

Nevertheless, during the last observing run [9], the *SkyHough* pipeline has experimented with the use of different statistics in each of the hierarchical steps. During the first step, candidates were ranked in




Figure 6.1: Example of the characteristic structures generated by continuous wave signal. Horizontal axis represent the parameter space region being evaluated; vertical axis represents the statistical significance of each template. Blue dots represent the obtained results during a first search over cleaned data; orange stars represent the obtained results after injecting a signal over the previously analyzed data.

terms of its critical ratio, as explained during chapter 3. This ensured a proper treatment of spectral disturbances, as the Hough method is one of the most robust strategy to deal with spectral disturbances. During the second step, once interesting regions of the parameter space have been identified, the analysis is performed using the *Powerflux* approach, which results in a more significant analysis thanks to the extra amount of information taken into account.

Hence, the universal statistic procedure could be implemented into the second hierarchical step of the *SkyHough* pipeline: Once the most significant candidate of a particular region has been identified, a parameter space sampling, completely analogous to the one presented in chapter 5, could be implemented to estimate the surrounding statistic. Then, the prescription given by [26] would be used to construct a conservative quantile and, in the end, a power upper limit would be set using (6.8). Finally, the upper limit could be propagated to any interesting quantity, such as h_0 or ϵ .

This formulation would join the inherent robustness of the *SkyHough* pipeline with an efficient way of computing upper limits, yielding a significant reduction of the computational cost of a continuous wave search.

6.3 Post Processing Reformulation

Finally, the last interesting improvement concerns the formulation of the post processing stage of the current pipeline. Nowadays, the coincidence and clustering windows have to be manually calibrated. Even though it yields a significant result, a different approach can be taken in order to cluster candidates in the parameter space.

Figure 6.1 represents the statistical imprint of an injection in the parameter space, as well as that of pure background noise. As clearly seen, the significance is not concentrated on a single template; rather, it forms a characteristic structure around the injection point.

The recognition of this kind of structures is a well-known problem which has lead to the development of interesting solutions within the framework of Machine Learning [29]. In particular, the classification of candidates into two classes (e.g. *noise* and *signal*) can be suitably done with the use of an Artificial Neural Network (ANN).

This type of algorithms use examples to characterize the role of different data features into the classification process. The calibration process, usually referred to as *training*, employs actual data samples to fit a set of free parameters of the model. Once the calibration is done, the classification process can be performed in a very efficient way thanks to the computational architecture of the algorithm.

The training of our model can be done by using several injections such as the ones shown in figure 6.1. The first step would be to *label* the candidates, depending on whether they are due to noise or a signal. To do so, we can construct a criterion by comparing the critical ratio of each template *before* and *after* injecting a signal into the data. Then, several neural network architectures would be trained, in order to look for the most accurate choice.

The algorithm is computationally feasible. It represents an improvement in the sense that the post processing would be calibrated using the information provided by the data itself, rather than manually choosing the coincidence and clustering windows. Moreover, once the training is done, the evaluation of an ANN is computationally inexpensive.

6.4 Summary

We presented three different works under development that may become a relevant contribution to the *SkyHough* pipeline after an initial development stage.

The τ -statistic is the most mature of them, as its coding is being tested and debugged to ensure its statistical correctness. It will improve the coherent step of the current pipeline, enhancing the detection of signals in a higher sensitivity depth. Its application will not produce a significant computational cost, as the procedure has to be done only *once* for each interesting frequency band into the set of SFTs.

The universal statistic will provide a an accurate estimation of the sensitivity of the search almost automatically after the determination of the most significant candidate. As a result, it could be used as a measurement on the good behavior of particular parameter space regions, easing the selection of useful segments if a further estimation of population based upper limits is needed.

Finally, the proposed reformulation of the post processing step in terms of an artificial neural network will introduce the machine learning techniques into the *SkyHough* pipeline, using the own structure of the parameter space to properly identify interesting candidates of an analysis.

The implementation of this developments will likely lead to a significant improvement of the *SkyHough* pipeline, further increasing the information extracted from the data.

Conclusion

The present work discusses the current state of continuous gravitational wave searches in terms of the current understanding of neutron stars as gravitational wave sources and the corresponding data analysis algorithms. Even though no continuous gravitational waves have been detected so far, the increasing sensitivity of the current interferometer detectors, together with the longer observing runs, leads the way towards a first detection. A direct implication of this fact is the establishment of more constraining upper limits as more observing runs become available.

Regarding computational capabilities, we discussed the main data analysis approaches, which are based on the segmentation of wide searches using semicoherent methods in order to identify interesting candidates which may be managed with a more powerful, coherent algorithm. The fact that there exists a wide variety of such algorithms eases the understanding of computational challenges and ensures an accurate coverage of any possible continuous wave source.

Among such algorithms, we find the *SkyHough* pipeline, one of the most noise-robust algorithms in the hunt for continuous waves. Its formulation mitigates the contribution of noisy sources into the analysis, preventing it to drift away due to a non-astrophysical measurement of the detector. The method is being continuously improved in order to keep its position at the forefront continuous gravitational wave searches.

We put the main focus on the statistical formulation of the pipeline. As discussed in chapter 5, we introduced two improvements into *SkyHough*:

Foremost, we implemented a sampling procedure to estimate the statistical properties of the background noise statistic. This is intended to take into account the actual background statistic, rather than an approximation which could overestimate the importance of a random fluctuation of the noise. The result was already shown during section 5.4 of the same chapter, obtaining a mitigation on the noise statistic with respect to the one used of a prescribed distribution.

Second, we used the achieved mitigation of noise contributions to combine data from multiple detectors into a single statistical measurement. Had we not implemented the parameter space sampling, this step would not be useful, as the noisy contributions from different detectors would increase the number of false candidates. The effectiveness of this new formulation was tested using the S6 Mock Data Challenge data set, which was introduced in chapter 4. We achieved to detect a greater number of signals with respect to the previous formulation of the *SkyHough* pipeline, obtaining an objective increase in the sensitivity of the new statistical formulation, as more signals with a *weaker* prominence with respect to the background were detected. The results were discussed during sections 5.4.2 and 5.4.2 of chapter 5.

Moreover, we experimented with the capabilities of a set of data analysis strategies in order to improve certain aspects of the current pipeline.

First, we developed a code to take into account spectral leakage in our searches using the τ -statistic. The code is being tested and debugged to ensure its proper behavior. A discussion on the statistical foundation of the algorithm was presented in section 6.1. It is expected that this idea refines the performance of the coherent step of the *SkyHough* pipeline, introduced in section 3.1.

Second, an alternative upper limit procedure, the universal statistic, was explored in terms of its portability to the *SkyHough* pipeline. Although it can not be directly applied to the standard Hough method, further improvements developed during the second observing run O2 could ease their application to the second hierarchical step of the pipeline. A qualitative discussion on this statement was presented during section 6.2. The computational implementation was already discussed by [26]. This procedure could yield preliminary results on the actual sensitivity of a search at the same time it is being performed. If we were to set upper limits by injecting an ensemble of signals, as exposed during section 2.4, we could use those results to evaluate the more suitable regions to perform subsequent analysis.

Third, we are developing a machine learning procedure to calibrate the post processing step in terms of the information provided by data. We expect this new formulation to be computationally inexpensive, as it will make use of an artificial neural network. Regarding its training process, it will be done by means of software injections of simulated continuous wave signals. We refer to section 6.3 to the corresponding discussion.

Every proposed improvement follows a simple idea: To obtain as much information as possible from the available data. The consistent use of this approach may result in the development of a well-behaved pipeline against noise disturbances, as the background behavior will be intrinsically taken into account by the data analysis procedure in itself.

Summarizing, we explored an interesting set of data analysis strategies to be applied into the *SkyHough* pipeline for the search of continuous gravitational waves. We expect to further develop these strategies in order to apply them to the forthcoming continuous wave data analysis within the LIGO and Virgo Scientific Collaboration.

As our pipelines are able to detect signals under more adverse conditions, so do improve the capabilities of gravitational wave interferometers, hopefully leading us towards a successful first detection of a continuous gravitational wave signal.

Bibliography

- [1] J. Aasi et al. Application of a Hough search for continuous gravitational waves on data from the fifth LIGO science run. *Class. Quantum Gravity*, 31(8), 2014.
- [2] J. Abadie et al. All-sky search for periodic gravitational waves in the full S5 LIGO data. *Phys. Rev. D*, 85(2):022001, 2012.
- [3] J. Abadie et al. Sensitivity Achieved by the LIGO and Virgo Gravitational Wave Detectors during LIGO’s Sixth and Virgo’s Second and Third Science Runs. *arXiv:1203.2674*, pages 1–11, 2012.
- [4] B. Abbott et al. First all-sky upper limits from LIGO on the strength of periodic gravitational waves using the Hough transform. *Phys. Rev. D*, 72(10):1–22, 2005.
- [5] B. P. Abbott et al. All-Sky LIGO Search for Periodic Gravitational Waves in the Early Fifth-Science-Run Data. *Phys. Rev. Lett.*, 102(March):111102, 2009.
- [6] B. P. Abbott et al. All-sky search for periodic gravitational waves in the O1 LIGO data. *Phys. Rev. D*, 96(6):1–35, 2017.
- [7] B. P. Abbott et al. First narrow-band search for continuous gravitational waves from known pulsars in advanced detector data. *Phys. Rev. D*, 96(12):1–20, 2017.
- [8] B. P. Abbott et al. Full Band All-sky Search for Periodic Gravitational Waves in the O1 LIGO Data. *Phys. Rev. D*, 97:102003, 2018.
- [9] B. P. Abbott et al. All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO O2 data. *arXiv:1903.01901v2*, pages 1–26, 2019.
- [10] B. P. Abbott et al. Searches for Continuous Gravitational Waves from 15 Supernova Remnants and Fomalhaut b with Advanced LIGO. *Astrophys. Journa*, 875(2):122, 2019.
- [11] N. Andersson. Gravitational waves from neutron stars. *Proc. Int. Astron. Union*, 5(H15):137–140, 2009.
- [12] G. Ashton. Neutron stars as continuous gravitational wave emitters. In *LIGO P1700388-v3 Neutron Stars Futur. Res.*, Bonn, 2017.
- [13] P. Astone, A. Colla, S. D’Antonio, S. Frasca, and C. Palomba. Method for all-sky searches of continuous gravitational wave signals using the frequency-Hough transform. *Phys. Rev. D - Part. Fields, Gravit. Cosmol.*, 90(4):1–24, 2014.
- [14] P. Astone, S. D’Antonio, S. Frasca, and C. Palomba. A method for detection of known sources of continuous gravitational wave signals in non-stationary data. *Class. Quantum Gravity*, 27(19), 2010.

- [15] G. S. Bisnovatyi-Kogan. The Neutron Star Population in the Galaxy. In B. Barbuy and A. Renzini, editors, *Proc. 149th Symp. Int. Astron. Union*, number December, page 379, Dordrecht, 1992. Kluwer Academic Publishers.
- [16] R. Biswas et al. Application of machine learning algorithms to the study of noise artifacts in gravitational-wave data. *arXiv:1303.6984v1*, pages 1–21, 2013.
- [17] R. D. Blandford and R. W. Romani. On the interpretation of pulsar braking indices. *Mon. Not. R. astr. Soc.*, 234:57–60, 1988.
- [18] P. R. Brady, T. Creighton, C. Cutler, and B. F. Schutz. Searching for periodic sources with LIGO. *Phys. Rev. D*, 57(4):2101–2116, 1998.
- [19] A. Bruce, M. A. Papa, and B. F. Schutz. Optimal strategies for sinusoidal signal detection. *Phys. Rev. D*, 66(10):1–18, 2002.
- [20] B. W. Carroll and D. A. Ostlie. *An Introduction to Modern Stellar Astrophysics*. Pearson Addison Wesley, San Francisco, CA, 2nd edition, 2007.
- [21] T. L. S. Collaboration. LIGO : The Laser Interferometer Gravitational-Wave Observatory. *arXiv:0711.3041v2*, 2009.
- [22] P. B. Covas et al. Identification and mitigation of narrow spectral artifacts that degrade searches for persistent gravitational waves in the first two observing runs of Advanced LIGO. *Phys. Rev. D*, 97(8):1–21, 2018.
- [23] P. B. Covas and A. M. Sintes. BinarySkyHough: a new method to search for continuous gravitational waves from unknown neutron stars in binary systems. *arXiv:1904.04873v1*, pages 1–18, 2019.
- [24] L. S. De La Jordana and A. M. Sintes. A chi2 veto for continuous gravitational wave searches. *Class. Quantum Gravity*, 25(18), 2008.
- [25] L. S. De La Jordana and The LIGO Scientific Collaboration and the Virgo Collaboration. Hierarchical Hough all-sky search for periodic gravitational waves in LIGO S5 data Hierarchical Hough all-sky search for periodic gravitational waves in LIGO S5 data. *J. Phys. Conf. Ser.*, 228:012004, 2010.
- [26] V. Dergachev. Novel universal statistic for computing upper limits in an ill-behaved background. *Phys. Rev. D - Part. Fields, Gravit. Cosmol.*, 87(6), 2013.
- [27] R. O. Duda and P. E. Hart. Use of the Hough transformation to detect lines and curves in pictures. *Commun. ACM*, 15(1):11–15, 1972.
- [28] T. D. Gebhard, N. Kilbertus, and I. Harry. Convolutional neural networks: a magic bullet for gravitational-wave detection? *arXiv:1904.08693v1*, pages 1–19, 2019.
- [29] A. Geron. *Hands-On Machine Learning with Scikit-Learn and TensorFlow*. O'REILLY, 2017.
- [30] GNU. GSL - GNU Scientific Library. *GNU Sci. Libr. - Free Softw. Found.*, 2019.
- [31] E. Goetz and K. Riles. An all-sky search algorithm for continuous gravitational waves from spinning neutron stars in binary systems. *Class. Quantum Gravity*, 28(215006), 2011.
- [32] T. Gold. Rotating Neutron Stars as the Origin of the Pulsating Radio Sources. *Nature*, 218(5143), 1968.
- [33] P. E. Hart. How the Hough Transform Was Invented. *IEEE Signal Process. Mag.*, 26(6)(November):18–22, 2009.

-
- [34] A. Hewsh, S. J. Bell, J. D. H. Pilkington, P. F. Scott, and R. A. Collins. Observation of Rapidly Pulsating Radio Source. *Nature*, 217:709–713, 1968.
 - [35] G. Hobbs, R. N. . Manchester, and L. Toomey. ATNF Pulsar Catalogue.
 - [36] M. M. Hohle, C. Ginski, J. G. Schmidt, and T. O. B. Schmidt. The companion candidate near Fomalhaut - a background neutron star? *Mon. Not. R. Astr. Soc.*, 448(1):376–389, 2015.
 - [37] C. J. Horowitz and K. Kadau. Breaking strain of neutron star crust and gravitational waves. *Phys. Rev. Lett.*, 102(19):1–4, 2009.
 - [38] P. V. C. Hough. Method and means for recognizing complex patterns. *US Pat. 3,069,654*, 21:225–231, 1962.
 - [39] <https://gwcenter.icrr.u-tokyo.ac.jp>. KAGRA.
 - [40] <https://www.caltech.edu/>. California Institute of Technology - LIGO.
 - [41] P. Jaranowski, A. Królak, and B. F. Schutz. Data analysis of gravitational-wave signals from spinning neutron stars: The signal and its detection. *Phys. Rev. D*, 58(6):063001, 1998.
 - [42] N. K. Johnson-Mcdaniel and B. J. Owen. Maximum elastic deformations of relativistic stars. *Phys. Rev.*, 88(4):1–20, 2013.
 - [43] D. I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. R. Astron. Soc.*, 331:103–2002, 2002.
 - [44] D. I. P. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon. Not. R. Astron. Soc.*, 331:203–220, 2002.
 - [45] E. Jones, T. Oliphant, P. Peterson, and E. Al. SciPy: Open Source Scientific Tools for Python. *Open Source*, 2001.
 - [46] B. Krishnan et al. Hough transform search for continuous gravitational waves. *Phys. Rev. D*, 70(8):1–22, 2004.
 - [47] B. Krishnan and A. M. Sintes. Hough search with improved sensitivity. *LIGO Doc. Control Cent.*, LIGO-T0701, 2007.
 - [48] J. M. Lattimer. The Nuclear Equation of State and Neutron Star Masses. *Annu. Rev. Nucl. Part. Sci.*, 62:458–515, 2012.
 - [49] L. Le Cam. An Approximation Theorem for the Poisson Binomial Distributino. *Pacific J. Math.*, 10(4):1181–1198, 1959.
 - [50] Ligo.org. LIGO Scientific Collaboration.
 - [51] D. R. Lorimer. Binary and Millisecond Pulsars Imprint. *Living Rev. Relativ.*, 8(11), 2008.
 - [52] A. G. Lyne, C. A. Jordan, C. M. Espinoza, B. W. Stappers, and P. Weltevrede. 45 years of rotation of the Crab pulsar. *Mon. Not. R. Astron. Soc.*, 446:857–864, 2015.
 - [53] M. Maggiore. *Gravitational waves. Volume 1: theory and experiments*, volume 1. Oxford University Press Inc., New York, first edition, 2008.
 - [54] S. Mahmoodifar and T. Strohmayer. Upper bounds on r -mode amplitudes from observations of low-mass x-ray binary neutron stars. *Astrophys. J.*, 773(2):140, 2013.
 - [55] R. N. Manchester. Millisecond Pulsars, their Evolution and Applications. *J. Astrophys. Astron.*, 38(3):1–18, 2017.

- [56] R. N. Manchester, G. B. Hobbs, A. Teoh, and M. Hobbs. The Australia Telescope National Facility Pulsar Catalogue. *Astron. J.*, 129(4):1993–2006, 2005.
- [57] C. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. W. H. Freeman and Company, San Francisco, 1st edition, 1973.
- [58] N. Nethercote and J. Seward. Valgrind : A Framework for Heavyweight Dynamic Binary Instrumentation. *ACM SIGPLAN Not. - Proc. 2007 PLDI Conf.*, 42(6):89–100, 2007.
- [59] M. Oliver, D. Keitel, and A. M. Sintes. The Adaptive Transient Hough method for long-duration gravitational wave transients. *Phys. Rev. D*, 99:104067, 2019.
- [60] C. Palomba, P. Astone, and S. Frasca. Adaptive Hough transform for the search of periodic sources. *Class. Quantum Gravity*, 22(18), 2005.
- [61] R. Prix. Gravitational waves from neutron stars. In W. Bejker, editor, *Neutron Stars Pulsars*, pages 137–140. Springer-Verlag, 2009.
- [62] R. Prix. Continuous Gravitational Waves from Spinning Neutron Stars. In *LIGO G1702407-v3 DPG-Frühjahrstagung*, 2018.
- [63] K. Riles. Recent searches for continuous gravitational waves. *Mod. Phys. Lett. A*, 32(39), 2017.
- [64] P. R. Saulson. *Fundamentals of interferometric gravitational wave detectors*. World Scientific, River Edge, NJ, 1st edition, 1994.
- [65] P. R. Saulson. If light waves are stretched by gravitational waves, how can we use light as a ruler to detect gravitational waves? *Am. J. Phys.*, 65(6):501–505, 1997.
- [66] I. I. Shapiro. Fourth Test of General Relativity. *Phys. Rev. Lett.*, 13(26):789, 1964.
- [67] A. M. Sintes and B. Krishnan. Improved Hough search for gravitational wave pulsars. *J. Phys. Conf. Ser.*, 32(1):206–211, 2006.
- [68] S. J. A. Test, S. Hypothesis, K. Poppenhaeger, K. Auchettl, and S. J. Wolk. A Test of the Neutron Star Hypothesis for Fomalhaut b A Test of the Neutron Star Hypothesis for Fomalhaut b. *Mon. Not. R. Astron. Soc*, pages 4018–4024, 2017.
- [69] D. Thain, T. Tannenbaum, and M. Livny. Distributed Computing in Practice : The Condor Experience. *Comput. Sci. Dep. Univ. Wisconsin-Madison*, (John Wiley & Sons, Ltd.), 2004.
- [70] The Continuous Waves Search Group. SFT Data Format Version 2 Specification. *LIGO-DCC*, (T040164-01-Z), 2004.
- [71] The LIGO Scientific and the Virgo Collaboration. All-sky search for periodic gravitational waves in LIGO S4 data. *Phys. Rev. D*, 77:1–38, 2009.
- [72] The LIGO Scientific Collaboration. Advanced LIGO. *Class. Quantum Gravity*, 32(7):74001, 2015.
- [73] The LIGO Scientific Collaboration. LIGO Algorithm Library - LALSuite <https://lscsoft.docs.ligo.org/lalsuite/>, 2018.
- [74] The LIGO Scientific Collaboration and the Virgo Collaboration. First search for gravitational waves from known pulsars with Advanced LIGO. *Astrophys. J.*, 839(12), 2017.
- [75] The LIGO Scientific Collaboration and the Virgo Collaboration. Narrow-band search for gravitational waves from known pulsars using the second LIGO observing run. *arXiv:1902.08442 [gr-qc]*, pages 1–20, 2019.

-
- [76] The LIGO Scientific Collaboration, the Virgo Collaboration, and A. Authors. Searches for Gravitational Waves from Known Pulsars at Two Harmonics in 2015-2017 LIGO Data. *arXiv:1902.08507*, 2019.
 - [77] The LIGO Scientific Collaboration and the Virgo Collaboration. Comprehensive all-sky search for periodic gravitational waves in the sixth science run LIGO data. *Phys. Rev. D*, 94:042002, 2016.
 - [78] M. Trias Cornellana. *Gravitational wave observation of compact binaries*. PhD thesis, University of Balearic Islands, 2010.
 - [79] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens. The LIGO Open Science Center, 2015.
 - [80] S. Walsh, M. Pitkin, M. Oliver, S. D. Antonio, V. Dergachev, A. Królak, P. Astone, and M. Bejger. Comparison of methods for the detection of gravitational waves from unknown neutron stars. *Phys. Rev. D*, 94(124010), 2016.
 - [81] D. H. D. West. Updating Mean and Variance Estimates : An Improved Method. *Commun. ACMS*, 22(9), 1979.
 - [82] D. R. Williams. Earth Fact Sheet, 2017.
 - [83] G. Woan, M. D. Pitkin, B. Haskell, D. I. Jones, and P. D. Lasky. Evidence for a minimum ellipticity in millisecond pulsars. *Astrophys. J. Lett.*, 863(2):L40, 2018.
 - [84] [Www.geo600.org/](http://www.geo600.org/). GEO.
 - [85] [Www.virgo-gw.eu/](http://www.virgo-gw.eu/). Virgo Scientific Collaboration.
 - [86] F. Zwicky and W. Baade. Cosmic Rays from Super-Novae. *Proc. Natl. Acad. Sci.*, 20(5):259–264, 1934.