UNIVERSITÀ
DEGLI STUDI
DI PADOVA

Department of Physics and Astronomy "Galileo Galilei"

Master thesis in Phisics of Data

Deep Learning methods to study the auxiliary channels in the Advanced Virgo Detector

Supervisor
Prof. Giacomo Ciani
External Supervisor
Prof. Massimiliano Razzano

Master Candidate
Luca Negri

Academic Year
2021/2022

Abstract

Gravitational waves are a prediction of Einstein's General theory of relativity and were first detected by the Advanced LIGO interferometers in 2015, almost 100 years after they were theorized. The detection of Gravitational Waves is affected by the background noise, which can be non-stationary and non-Gaussian. Short-duration noise bursts, also called "glitches" can be quite detrimental to the search, since they can affect data quality and mimic the gravitational waves signal itself. To assert the quality of the measurements and the origins of the glitches, thousands of sensors monitor the state of the gravitational wave detectors continuously, but the large amount of data and the complex nature of the couplings between channels render any kind of analysis quite challenging. Deep Learning offers an opportunity to efficiently analyze and handle large amounts of data related to the Auxiliary Channels. This thesis aimed to apply Deep Learning tools to the study of the Auxiliary Channels to better characterize and gain further insight into the noise of the Advanced Virgo detector. In particular, Variational AutoEncoders (VAE) can help in the task of reducing the dimensionality of the data without losing any relevant information, by projecting the samples on a lower dimensional manifold, called the latent space, on which the analysis can be run instead.

In this thesis we exploit the latent space representation of a multi-channel Variational Autoencoder trained on Auxiliary Channels spectrograms, coupled with a clustering algorithm to find patterns in this complex data landscape.

The main focus was to study the response of the interferometer's mirror suspensions, called Superattenuators, to small-scale seismic events and to explore the correlation between data quality and low-frequency glitches in the main channel. This is one of the first projects of this kind that analyzes gravitational waves detectors data and should be taken as a proof of concept of how machine learning can be integrated into this field. The results are promising for the characterization of the performances of Superattenuators during small earthquakes and in identifying the causes of some glitches, which is a step towards improving the detector sensitivity in the low-frequency domain.

Contents

Introduction

Gravitational waves are one of the newest and most active fields in astrophysics, as they provide a new window to look at yet to be studied parts of the universe. These incredibly small perturbations induce changes in the length of the order $\mathcal{O}(10^{-18})$m over 1 km, making their detection one of the last predictions of general relativity to be experimentally confirmed. The gravitational wave detectors of the LIGO-Virgo-KAGRA (LVK) collaboration are in a constant fight to reduce as much as possible the noises that affect their measurements, by constantly improving the capabilities of these instruments and mitigating known noise sources. To investigate these causes, the scientists are helped by the auxiliary channels of the interferometers, generated by many sensors that monitor the subsystems and the surrounding environment of the instrument at all times. The job of monitoring the state of the interferometer and its subsystem is a hard task: non-linear couplings can make it difficult for analytical algorithms to find correlations between the subsystems and the detector output, while the vast amount of data produced by the auxiliary channels renders any kind of manual analysis impractical. But these immense datasets can be an invaluable resource for the training of machine learning algorithms.

This thesis work places itself in this new emerging field, presenting one of the first works that use machine learning to characterize the response of the seismic isolation systems of the Advanced Virgo interferometer, the Superattenuators, to short bursts of seismic activity.

More and more research groups are trying to develop novel methods of data analysis to tackle these problems, and many of these are using a machine learning driven approach. Machine learning can provide statistical analysis of arbitrary complex phenomena, given enough training data, and the huge amount of data acquired during the years of operation of the LVK collaboration could allow the training of deep learning models to try and understand better these complicated non-linear behaviors.

Gravitational wave signals tend to evolve from lower to higher frequencies. To measure these signals for longer times, it is necessary to improve the detector performance at low frequencies. For these reasons, the mirror suspension system, the SuperAttenuator, was chosen as the study objective of the thesis since seismic noise is the main limiting factor in the lower parts

of the spectrum. Not many previous works exist on this subject, so the main goal of this project is to present a proof of concept of how machine learning could improve our knowledge of these complex systems. For these reasons the scope of the work was kept quite narrow, focusing on characterizing the position of the upper part of the Superattenuator, called the Filter 0, and how it responds to seismic activity.

The actual thesis work was to develop an algorithm capable of finding patterns in the responses of the Superattenuators to seismic bursts, by first building a generative model for the dataset, capable of reducing the dimensionality of the problem, and then by running a clustering algorithm in the now less complex data. By examining those clusters some insight could be gained into how different Superattenuators perform, what different noises affect the suspensions and what perturbations have the greatest effects on the detector output.

Chapter 1 will introduce the main theoretical framework behind gravitational waves. It will discuss how these perturbations emerge from solving Einstein's field equations and what is the current state of gravitational wave science.
Chapter 2 will describe the Advanced Virgo interferometer, giving an overview of the science behind the detector and how these principles are implemented in the field.
Chapter 3 will give an overview of the machine learning tools used in the project, and how these novel analysis methods have already been successfully used in gravitational waves astronomy.
Chapter 4 describes how the theoretical tool previously presented are implemented in the algorithm to build the dataset of responses of the Superattenuators, with an analysis to better understand the biases that the algorithms might encounter.
Finally, chapter 5 will explore the generative model performances and the results of the clusters analysis, to look for insights on how algorithms like this could help improve the detector performances in the future.

Chapter 1

Gravitational Waves

This Chapter will give an overview of the science behind gravitational waves. Sec. 1.1 will serve as an introduction to general relativity and how it allows the propagation of gravitational waves will be discussed in sec. 1.2 . A general look at the physical effects of gravitational waves will be given in sec. 1.3 while their generation will be explored in sec. 1.4. Finally, an overview of the experimental detection of gravitational waves and the current state of this field of physics will be given in sec. 1.5.

1.1 General Relativity

Einstein's general theory of relativity represents our most comprehensive description of gravity [1]. It describes how the distribution of matter and energy can influence the curvature of spacetime. In general relativity, the energy density at a certain point is described by the stress-energy tensor $T_{\mu\nu}$, while the spacetime metric is defined as $g_{\mu\nu}$, and the interaction between the two is mediated by Einstein's field equations

$$R_{\mu\nu} - \frac{1}{2} g_{\mu\nu} R = \frac{8\pi G}{c^4} T_{\mu\nu} \tag{1.1}$$

The $R_{\mu\nu} = R^{\gamma}_{\mu\gamma\nu}$ and $R = g^{\mu\nu} R_{\mu\nu}$ therms are respectively the Ricci Tensor and the Ricci scalar which are contraptions of the more complex Riemann tensor $R^{\mu}_{\nu\tau\rho}$, which encodes the properties of the curvature of spacetime. This is obtained by a non-linear combination of the Christoffel symbols and their derivatives, which is a non-tensorial object that encoded the differences between flat and curved spacetime [1]

$$\Gamma^{\mu}_{\nu\epsilon} = \frac{1}{2} g^{\mu\lambda} (\partial_{\epsilon} g_{\lambda\nu} + \partial_{\nu} g_{\epsilon\lambda} - \partial_{\lambda} g_{\nu\epsilon}) \tag{1.2}$$

From this, the Riemann tensor can be obtained:

$$R^{\mu}_{\nu\tau\rho} = \partial_{\tau} \Gamma^{\mu}_{\nu\rho} - \partial_{\rho} \Gamma^{\mu}_{\nu\tau} + \Gamma^{\mu}_{\lambda\tau} \Gamma^{\lambda}_{\nu\rho} + \Gamma^{\mu}_{\lambda\rho} \Gamma^{\lambda}_{\nu\tau} \tag{1.3}$$

1.2 Linear Solutions to the field equations

Solutions to Einstein's field equation arise from the 10 non-linear differential equations, but the complexity can be drastically reduced by making some assumptions. Is first assumed that the spacetime studied will be a small perturbation $\|h_{\mu\nu}\| \ll 1$ of flat spacetime represented by the Minkowsky metric $\eta_{\nu\nu}$. The total metric will be

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \tag{1.4}$$

Solving the field equation for $g_{\mu\nu}$ and approximating at first order in $h_{\mu\nu}$ [2]:

$$\Box\overline{h}_{\mu\nu} + \partial^\rho\partial^\sigma\overline{h}_{\rho\sigma}\eta_{\mu\nu} - \partial^\rho\partial_\nu\overline{h}_{\mu\nu} - \partial^\rho\partial_\mu\overline{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \tag{1.5}$$

where $\overline{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$ and h is the trace of $h_{\mu\nu}$. The presence of the d'Alambertian operator $\Box$ hints towards a wave-like solution for $h_{\mu\nu}$, but the other terms render this intuition hard to apply. At this stage, there exist multiple possible solutions for $h_{\mu\nu}$ that describe the same spacetime. To find a meaningful form for the metric it is necessary to set a specific frame of reference, where the equations will assume a simpler form. This ambiguity in the equations is called gauge freedom, and the right choice of gauge will lead to a more understandable solution. In this case, the harmonic gauge (sometimes called Lorentz gauge) is the one that will make the equations assume their simplest form. By applying a small coordinate change ξ^μ of order $\mathcal{O}(h)$ $\widetilde{x}^\mu = x^\mu + \xi^\mu(x)$ the new metric approximated at $\mathcal{O}(h)$ will take form

$$\widetilde{h}_{\mu\nu} = h_{\mu\nu} + \partial_\nu\xi_\mu + \partial\mu\xi_\nu \tag{1.6}$$

It can be shown [1] that it is always possible to find a change of coordinate such that $\Box\xi^\mu = 0$, which implies $\partial^\mu\widetilde{\overline{h}}_{\mu\nu} = 0$. This condition fixes 4 of the 10 starting degrees of freedom and, with this new Gauge, Einstein's field equations collapse to

$$\Box\overline{h}_{\mu\nu} = -\frac{16\pi G}{c^4}T_{\mu\nu} \tag{1.7}$$

To study how gravitational waves propagate, the equations will be solved in the vacuum case $\Box\overline{h}_{\mu\nu} = 0$. The more complex solutions where $T_{\mu\nu} \neq 0$ will be looked at in sec. 1.4. As mentioned previously, the d'Alambertian hints towards plane waves as solutions:

$$\overline{h}_{\mu\nu} = H_{\mu\nu}e^{ik_\mu k^\mu} \tag{1.8}$$

where $H_{\mu\nu}$ is the polarization matrix and $k_\mu = (\omega/c, \mathbf{k})$ is the wave vector. For this to be a valid solution, the condition $\omega/c = \|\mathbf{k}\|$ must hold, implying that $\mathbf{k}$ is a null vector. Null vectors are typical of photons and other force-carrying particles, meaning that gravitational waves propagate at the speed of light [1]. The last step needed to obtain the full equation for

gravitational waves is to find a meaningful form for $H_{\mu\nu}$ and this can be accomplished by fixing the remaining degrees of freedom.

Since multiple ξ_μ exists for which $\Box\xi_\mu = 0$, it can be shown that it is possible to set $h^{0i} = 0$ and still be in the Lorentz Gauge, so in this frame of reference the gravitational wave has no time components. One last degree of freedom can be fixed by setting the trace of the matrix $h^\mu_\mu = 0$. The full set of equations that fixes the degrees of freedom are

$$\begin{cases} \partial^i h_{ij} = 0 \\ h^{0\mu} = 0 \\ h^i_i = 0 \end{cases}, \tag{1.9}$$

The first condition implies that along the direction of motion of the wave, the metric remains constant making gravitational waves transverse. This is the reason why this set of conditions is commonly referred to as the *Traceless-Trensverse* gauge, or TT-Gauge. In equation 1.9 the first condition is the harmonic gauge, which fixes 4 degrees of freedom, while the time independence and the traceless conditions fix the other 4 . This means that only 2 degrees of freedom are remaining. In conclusion, the real part of a wave moving in the z direction will have the equation[2]:

$$h^{TT}_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_+ & h_\times & 0 \\ 0 & h_\times & -h_+ & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \cos[\omega(t - z/c)] \tag{1.10}$$

In this linear approximation, gravitational waves can be completely characterized by the wave vector k^μ and the two remaining degrees of freedom, encapsulated in the polarization matrix by the variables h_+ (h-plus) and $h_\times$ (h-cross). The name derives from the fact that these variables affect spacetime at orientations of $\pi/4$ from each other. This is a key difference with electromagnetic waves, where the possible polarizations are oriented at $\pi/2$. An illustration of how these polarizations affect spacetime is seen in fig. 1.1

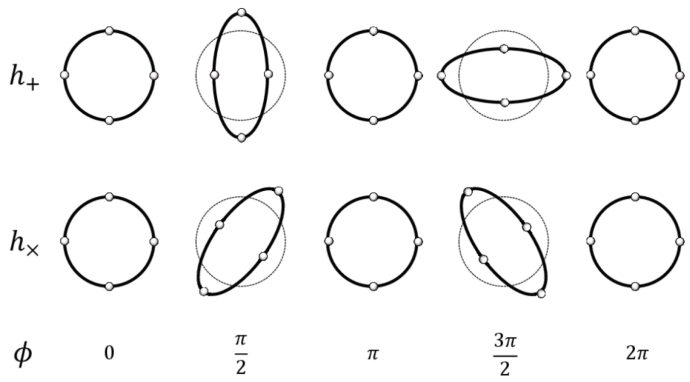

Figure 1.1: The possible polarization of gravitational waves and how they perturb spacetime [3]

1.3 Effects of gravitational waves

The change of the local metric induced by a gravitational wave will inevitably modify the path that geodesics will take through spacetime, implying that free-falling objects will experience a deviation from their trajectories. In general, the evolution of the proper distance between two free-falling objects L_μ will have the equation:

$$\frac{D^2 L^\mu}{D\tau^2} = R^\mu_{\nu\rho\sigma} \frac{dx^\nu}{d\tau} \frac{dx^\sigma}{d\tau} L^\rho \tag{1.11}$$

This is commonly referred to as the geodesic deviation equation. By fixing L_μ as the distance before the presence of a gravitational wave, the change δL_i of this quantity is obtained by solving equation 1.11 in linearized gravity [4]

$$\delta L_i = \frac{1}{2} h^{TT}_{\mu\nu} L^i_0 \tag{1.12}$$

which means that the change in proper length is proportional to the initial distance L^i_0 and the gravitational wave itself. If L^i_0 is oriented in the x direction, while the gravitational wave is going in the z direction with a plus polarization, the equation for the change in the distance will be

$$\mathbf{s} = L + \frac{Lh}{2} \cos(\omega t + \phi) \tag{1.13}$$

As will be discussed in Chapter 2, the measurement of this quantity is what detectors use to infer the presence of gravitational waves.

1.4 Sources of gravitational waves

1.4.1 Generation of Gravitational waves

To find how gravitational waves can be produced, it is necessary to find solutions to equation 1.7 where the Stress Energy tensor is different from 0. The generic solution of this equation will be found using the Green function G [2]

$$\overline{h}_{\mu\nu} = -\frac{16\pi G}{c^4}\int d^4x' G(x-x')T_{\mu\nu}(x') \tag{1.14}$$

To find analytical solutions a few assumptions must be made:

- The system should still operate under the linearized gravity assumption used so far (only Newtonian interactions between objects). Meaning that the equation will have a simple solution in the TT gauge.
- The system has a size L that is much smaller than the wavelength of the gravitational wave λ_{GW}. This implies that the objects involved move at speeds $v \ll c$. Under this assumption, the system can be expanded in $\frac{v}{c}$, and therm with order greater than $\mathcal{O}(\frac{v}{c})$ can be ignored.
- The solution is calculated for a point at distance r from the source, where r is much greater than L so that the influence of the single components of the system cannot be fully resolved and the effect of the system is considered as a whole.

All of these conditions can be summarized into

$$L \ll \lambda_{GW} \ll r \tag{1.15}$$

And the general solution at order $\mathcal{O}(\frac{1}{r})$, calculated at the retarded time $t_r = t - r/c$, will be [4]

$$\overline{h}_{\mu\nu}(x,t) = \frac{4G}{c^4}\frac{1}{r}\int d^3x' \left[T_{\mu\nu}(x',t_r) - \frac{1}{c}\dot{T}_{\mu\nu}(x',t_r)x' + \frac{1}{2c^2}\ddot{T}_{\mu\nu}(x',t_r)x^2\right] \tag{1.16}$$

The 3 terms of this equation correspond to the contributions of the monopole, dipole and quadrupole moments of $T_{\mu\nu}$ to the gravitational wave. Since, under these assumptions, the conservation law $\partial_\nu T^{\mu\nu} = 0$ must hold, the total mass and momentum of the system cannot change in time: this means that the monopole and dipole terms will stay constant. Their contribution to the gravitational wave will then be constant as well, meaning that it can be set arbitrarily to 0. Only quadrupole and higher moments of the system will contribute to the radiation of gravitational waves. The final solution approximated at first order can be written as [2]

$$\overline{h}_{ij}(t,r) = \frac{2G}{c^4}\frac{1}{r}\ddot{Q}_{ij}^{TT} \tag{1.17}$$

$$Q_{ij} = \frac{1}{c^2}\int d^3x' T_{00}(x', t_r) x'_i x'_j \tag{1.18}$$

The $\frac{2G}{c^4} \sim \mathcal{O}(10^{-44})\frac{s^2}{\text{mkg}}$ therm in equation 1.17 implies that only systems with a great quadrupole moment will have the ability to produce gravitational waves measurable by our instruments. For this reason, the sources of gravitational waves that are targeted by modern detectors have to be the most violent and energetic events in astrophysics. The rest of this section is dedicated to illustrating a list of these astrophysical sources

1.4.2 Coalescing binaries

Coalescing binaries are the only source of gravitational waves that have been directly detected so far [5]. These systems are binary stars where the components are both compact objects (either black holes or neutron stars, both of which have been observed emitting gravitational waves [6]). During the O3 observing run, the LIGO-Virgo collaboration has detected binary system mergers regularly, with the rate of about once a week [7]. Binary star systems are quite abundant in our universe, which makes the emergence of compact binaries even more probable. By calculating the quadrupole of the binary system while considering the approximations introduced in equation 1.15, the equation for the polarization of the gravitational waves becomes [1]

$$\begin{aligned} h_+(r,\theta,t) &= \tfrac{4}{r}\left(\tfrac{GM_c}{c^2}\right)^{5/8}\left(\tfrac{\pi f_{gw}}{c}\right)^{2/3}\tfrac{1+\cos^2\theta}{2}\cos(2\pi f_{gw}t + 2\phi) \\ h_\times(r,\theta,t) &= \tfrac{4}{r}\left(\tfrac{GM_c}{c^2}\right)^{5/8}\left(\tfrac{\pi f_{gw}}{c}\right)^{2/3}\cos^2\theta\sin(2\pi f_{gw}t + 2\phi) \end{aligned} \tag{1.19}$$

Where $f_{gw} = 2f_{orb}$ is the frequency of the gravitational wave, θ is the inclination of the orbital plane of the source with respect to the observer and M_c is the *chirp mass*:

$$M_c = \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}} \tag{1.20}$$

Compact binary systems tend to lose energy and angular momentum through the radiation of gravitational waves, which makes the radius of the orbit smaller. Smaller orbits mean faster and closer objects, leading to a higher quadrupole moment of the system and more gravitational waves being emitted. This will start a positive feedback loop for the radiative process that continues until the two bodies collide and coalesce into a single object. By defining τ as the time until coalescence and by approximating to circular orbits, by calculating the radiated power it is possible to write the equation that defines the evolution of the system as:

$$
\begin{aligned}
h_+(r,\theta,t) &= \frac{1}{r}\left(\frac{GM_c}{c^2}\right)^{5/4}\left(\frac{5}{c\tau}\right)^{1/4}\frac{1+\cos^2\theta}{2}\cos(\phi(\tau)) \\
h_\times(r,\theta,t) &= \frac{1}{r}\left(\frac{GM_c}{c^2}\right)^{5/4}\left(\frac{5}{c\tau}\right)^{1/4}\cos^2\theta\sin(\phi(\tau)) \\
\phi(\tau) &= -2\left(\frac{c^3}{5GM_c}\right)^{5/8}\tau^{5/8}+\phi_0
\end{aligned}
\tag{1.21}
$$

This equation is a valid approximation of the evolution of the system right until moments before coalescence, where the slow-motion assumptions fail. One interesting note is that the evolution of the system is entirely determined by the chirp mass, which is the most accurate quantity measured from coalescing binaries [8]. For bodies of chirp mass $\gtrsim 1M_\odot$ right before coalescence, where the energy radiated is at a maximum, the gravitational waves emitted reach frequencies in the range 10-1000 Hz, which is the range where most current detectors operate. The rise in frequency of the system until the final merger is called chirping and it is shown in fig. 1.2

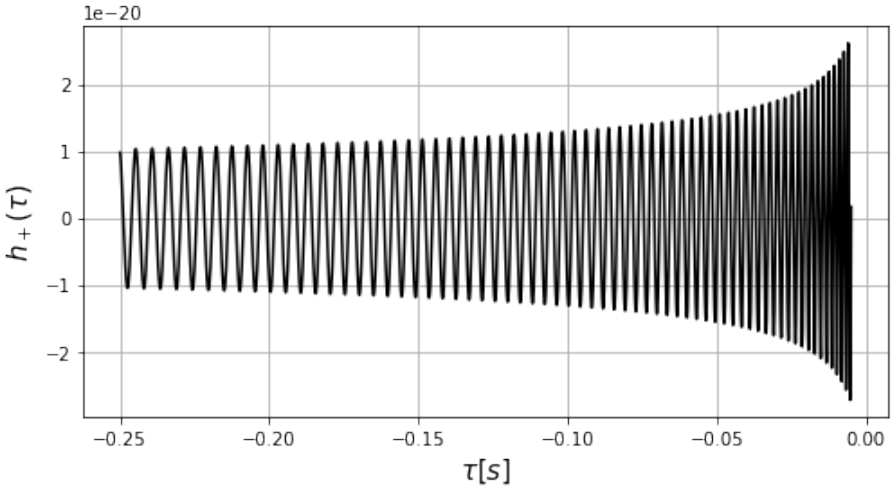

Figure 1.2: The Evolution of $h_+(\tau)$ for a coalescing neutron star binary, where both components have $2M_\odot$, starting at 0.25 second before coalescence. The simulation is run until the assumptions taken to write the eq.s 1.21 break down.

1.4.3 Supernovae

Supernovae are one of the most violent and energetic kinds of events that can happen in astrophysics. They are the endpoint of the life of the most massive classes of stars. The star collapses under its own weight and the outer material falls on the core at close to the speed of light, generating a violent explosion. The energy is released mostly through neutrinos, but there are significant electromagnetic and gravitational components [1]. The latter originates when the collapse is not spherically symmetric, meaning that a quadrupole moment can emerge. These events have not been detected through gravitational waves yet, but a lot of work has been done

to search for these events [9]. They are thought to be quite rare and are quite hard to model [10].

1.4.4 Rotating Neutron stars

Together with transient events, modern detectors are also trying to find continuous sources of gravitational waves. The most promising candidates in this category are the emission caused by rotating asymmetric neutron stars [11]. Neutron stars are the remnants of supernovae, that originate whenever the mass of the core of the star was $< 2M_{\odot}$. These objects are really small in size and present really high densities. Due to the conservation of angular momentum from when the star was larger in volume, these objects tend to spin quite fast: if some asymmetry is present on their surface, a quadrupole moment can emerge, and gravitational waves can be emitted. If a neutron star is rotating with frequency ω around the z axis of a star with I_x, I_y and I_z as moments of inertia, an observer at distance r and for which the source is inclined by an angle θ will measure the incoming gravitational waves as:

$$\begin{aligned} h_+(t) &= \frac{4G\omega^2\epsilon I_z}{c^4 r}\frac{1+\cos^2(\theta)}{2}\cos(2\omega t) \\ h_\times(t) &= \frac{4G\omega^2\epsilon I_z}{c^4 r}\cos(\theta)\sin(2\omega t) \end{aligned} \tag{1.22}$$

Where $\epsilon = (I_x - I_y)/I_z$.

The emissions from these sources tend to be much quieter than the ones produced by transient events, but their prolonged nature raises the SNR for the whole time they are in the detector range. So far none of these sources were found, and only upper bounds in their emission were put [12].

1.4.5 Stochastic background

The gravitational stochastic background is thought to have 2 major components. The first is of astrophysical origin and is composed of many unresolved sources, like inspiralling binaries far from coalescence, rotating pulsars and far away core-collapse supernovae [13]. The other is of cosmological origins, coming from the chaos of the times right after the big bang. These signals are predicted to affect the lowest part of the frequency spectrum, well below the capabilities of current detectors, but are the target for searches for future ones like the LISA observatory [14] and Einstein Telescope [13].

1.5 Observation of gravitational waves from ground-based detectors

Since 2015, the second generation of gravitational waves observatories has been looking for signals. The interferometers of the LVK collaboration were upgraded to Advanced LIGO [15]

and Advanced Virgo [16], and since they were put online they have confirmed the detection of 90 gravitational waves transient events all coming from the coalescence of compact binaries. These observations were divided into three observing runs and are presented together with the planned future ones, in fig. 1.3. A brief description of these runs will now be presented.

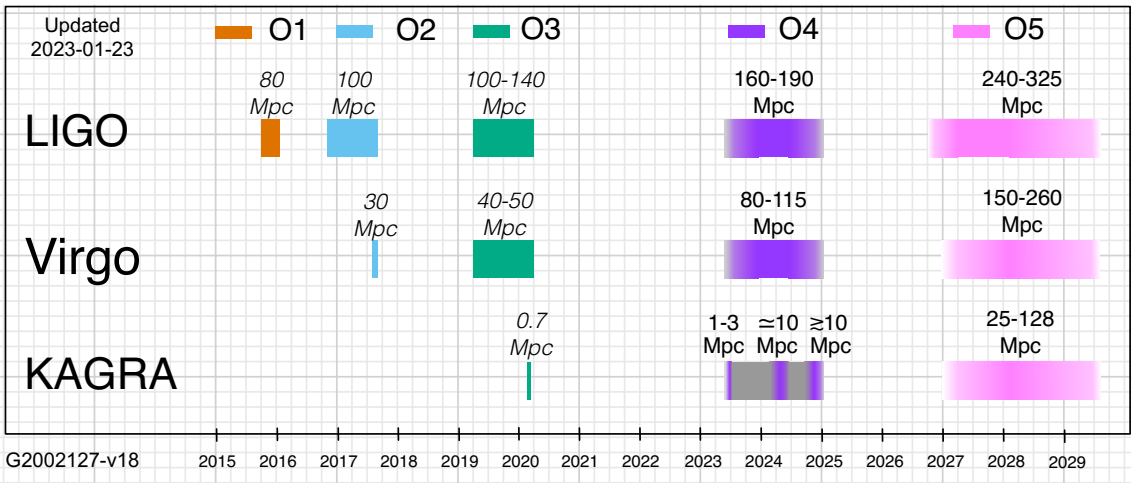

Figure 1.3: Timeline of the different observing runs of the LVK collaboration, both past and future prospects are present. [17]

- **O1**: from 12 September 2015 to 19 January 2016, the first observing run yielded the first-ever detection of gravitational waves with GW150914. In total 3 events have been detected by the network during this period. [8]

- **O2**: from 30 November 2016 to 25 August 2017, this run detected the first Binary neutron star merger GW170817, reported in fig. 1.4, that also had an electromagnetic observation follow-up. Also, the first event that had been observed by the whole three-detector network GW170814, greatly improving the sky triangularization. A total of eight events have been identified in this run.[8]

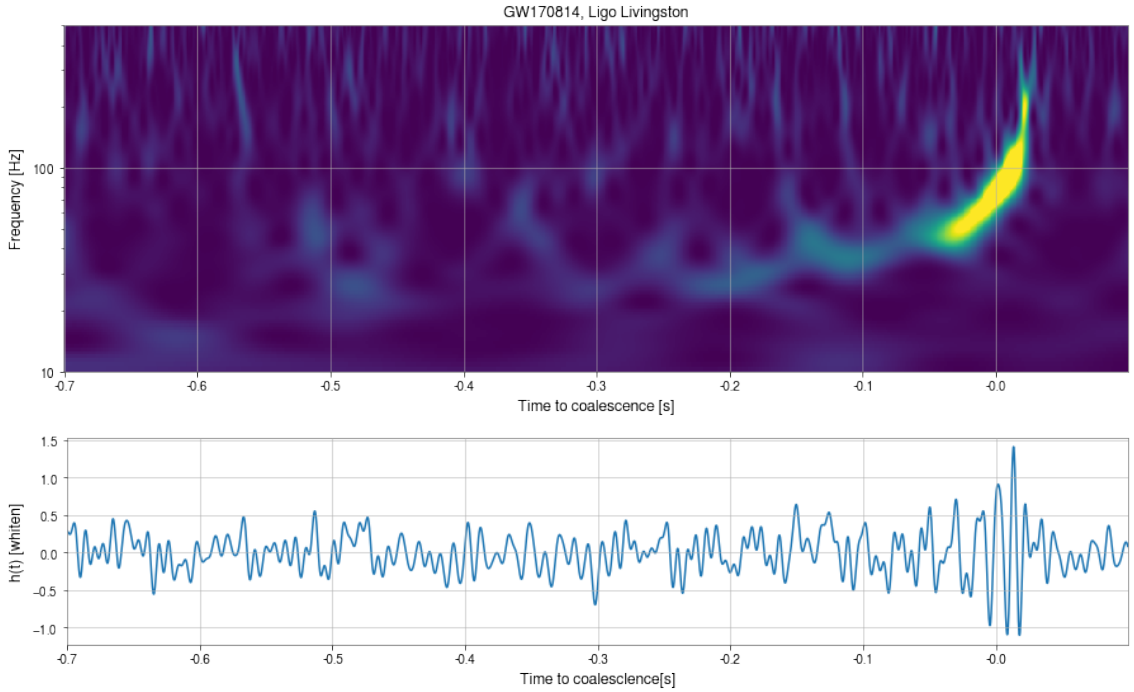

Figure 1.4: Above the Q-transform of the GW170814 event as measured by the LIGO Livingston interferometer around the time of coalescence. Below there is a whitened and low-pass fileterd version of the strain channel $h(t)$ for the same event. This event is the first three interferometer detection, which made it possible to have the best sky localization at the time

- **O3a** and **O3b**: the third observing run was divided into two halves, the first started April 1 and ended October 1 2019, identifying 39 events, the second started 1 November 2019, and ended 27 March 2020 identifying 35 events, signaling how the improvements in the precision of the detector yielded more and more possible candidates. Between these, GW190521 is the most massive system of binary black holes ever observed, with masses of at least one starting component of the system being an intermediate-mass black hole, the first ever observed. Also with GW190412, the network observed a highly asymmetrical system, which lead to the measurement of higher harmonics of the gravitational wave.[6], [7]

- **O4** and **O5**: These are the observing runs planned for the future. During these the network will be joined by the now fully capable KAGRA observatory in Japan, improving the sky localization and the total uptime of the network. Frequency-dependent light squeezing will be implemented for the whole network [18] and for Virgo in particular, heavier mirrors ($\sim$ 100 Kg) and the implementation of the Signal Recycling cavity are in the works [18]. These improvements will extend the range of possible detection by a factor of 2-3.

An overview of all the events ever witnessed by the LVK collaboration together with other electromagnetic observations is summarized in fig. 1.5

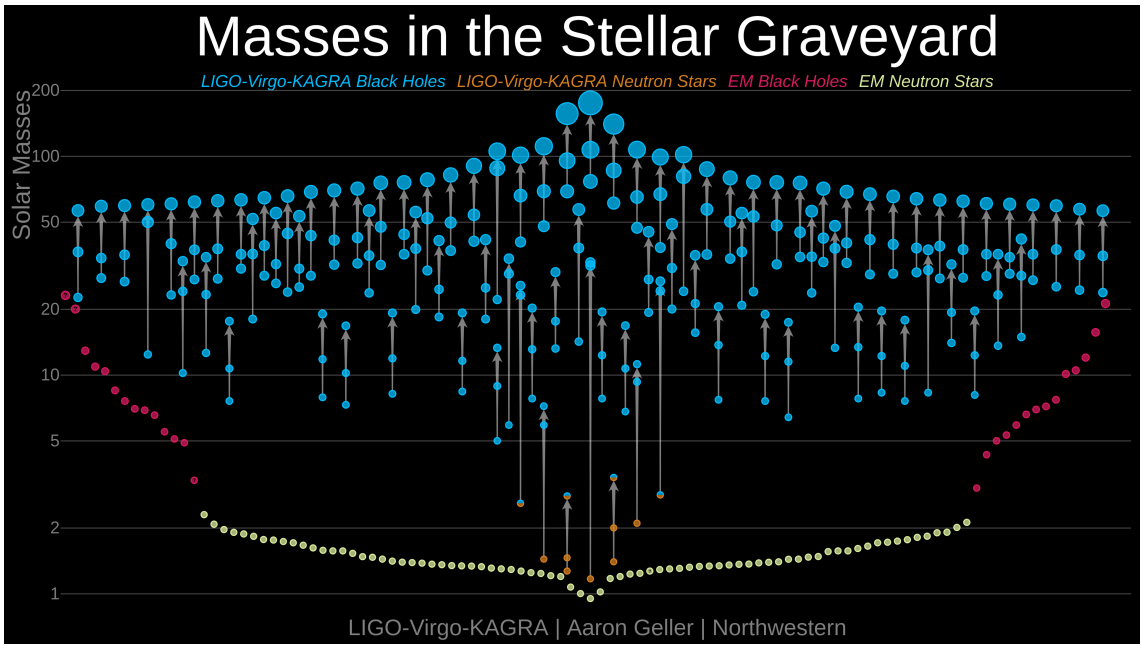

Figure 1.5: All of the known stellar remnants ever measured, with the EM observation at the bottom and the LVK collaboration observation at the top [19]

Summary

In this chapter, solutions to Einstein's field equation in a vacuum were proven to allow the propagation of gravitational waves. Their effects on free-falling masses and the methods through which they could originate were discussed next, by giving a list of the astrophysical sources that are currently the target of gravitational wave searches. Finally, an overview of the status of the current experimental observations of gravitational waves was given.

Chapter 2

Advanced Virgo

This chapter will give an overview of how gravitational waves are detected with the Advanced Virgo Interferometer, and the challenges that the instrument faces in achieving this task. In sec. 2.1 the working principles behind Michelson interferometers will be described, and sec. 2.1.2 will expand on how Fabry-Perot Cavities can improve their performances. In sec. 2.2 the layout of the Advanced Virgo interferometer will be presented and an overview on the mirrors suspensions, the Superattenuators, will be given in sec. 2.3, since these will be the focus of the thesis. Then, a general description of the noise affecting the detector is given in sec. 2.4, regarding static noise, and in sec. 2.5. Finally, the steps needed to get the interferometer to its full capabilities are described in sec. 2.6.

2.1 Michelson interferometers

Michelson interferometers have been so far the only instruments capable of detecting gravitational waves. This is due to their exceptional ability to measure really small length changes over long distances. In fig. 2.1 a schematic of a simple Michelson interferometer is pictured, and a brief description of its working principles will be presented, with a focus on why they are the instrument of choice when trying to detect gravitational waves.

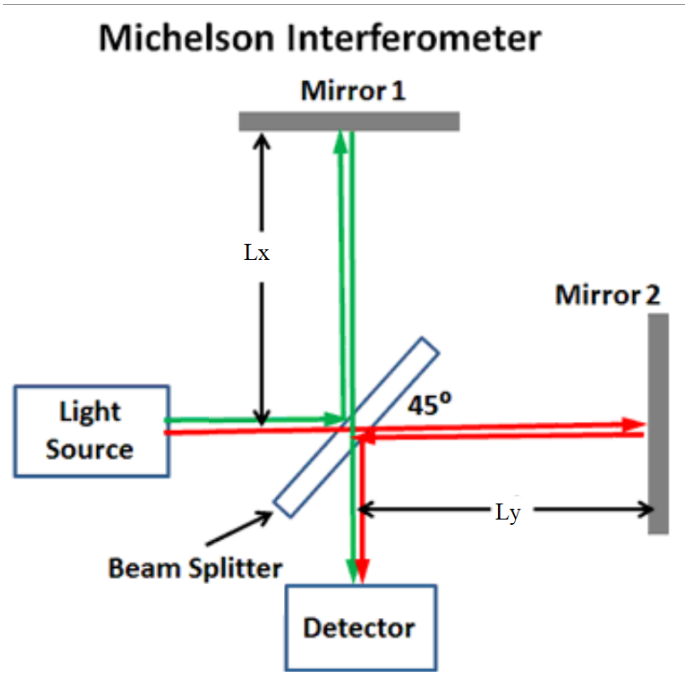

Figure 2.1: Simple Michelson interferometer layout[20]

The Michelson interferometer works by sending a beam of monochromatic light, emitted by a laser, to a semi-transparent mirror (the Beam-Splitter) that has a 50/50 chance of either transmitting or reflecting the incoming light [1]. If the Beam Splitter (BS) is oriented at $45°$ with respect to the incoming light, two perpendicular beams will form. After traveling for their designed distance, the beams get reflected back towards the BS, where they will recombine and interfere with each other. The light exiting from the BS in the opposite direction of the reflected beam will be sent to a photodetector to measure its intensity. At the beam splitter, the two returning beams will have the same wavelength, but a different phase, since this quantity is dependent on the length of the path that each beam took. The power measured by the photodetector will depend on this phase difference and if k_L is the wave vector of the laser light and the length of the two paths is L_x and L_y, the electric field at the output port will be:

$$E_{out}^2 \propto E_0^2 \sin^2[k_L(L_y - L_x)] \tag{2.1}$$

2.1.1 Gravitational waves detection with Michelson interferometers

If the two mirrors at the ends act as free-falling masses in the direction of the incoming beam, then the measured output can be affected by the passage of gravitational waves. To have free-falling objects on earth, these mirrors need to be suspended from the ground and left free to move. The suspension system of Advanced Virgo that makes this possible will be the focus of sec. 2.3. The free-falling property holds only for short timescales since the mirror movement

cannot be completely free, as to avoid the mirror wandering outside the beam reach. Current interferometers are focused on studying just frequencies above 10 Hz, so whenever this restoring force acts on longer timescales, this will not interfere with the free-falling property.

By taking as an example a h_+ polarized gravitational wave moving along the z direction with amplitude h_0, by integrating the path of a photon along the x arm of the interferometer in the TT-gauge is possible to obtain the phase shift at the beam-splitter induced by the gravitational wave [1]:

$$\Delta\phi_x(t) = h_0 k_L L \mathrm{sinc}(\omega_{gw} L/c) \cos[\omega_{gw}(t - L/c)] \tag{2.2}$$

Meanwhile, the result along the y arm will be $\Delta\phi_y = -\Delta\phi_x$ and the total phase difference between the beams will be $\Delta\phi_{mich} = \Delta\phi_x - \Delta\phi_y = 2\Delta\phi_x$. The sinc function arises from the fact that the effect of the gravitational wave is not constant during the flight of the photon. If $L > \lambda_{gw}/2$ the gravitational wave effect will change direction and intensity before the photon arrives at the BS, partially canceling its effect, and whenever $L = n\lambda_{gw}$ this will result in a displacement of 0. The effect of the gravitational wave can be considered constant only for $L \ll \lambda_{gw}$, and in this regime, the power observed at the photodetector will be:

$$P = P_0[1 - \cos(2\phi_0 + \Delta\phi_{mich})] \tag{2.3}$$

The initial phase ϕ_0 defines the working point of the interferometer, and its choice is fundamental to have a capable instrument. The strategy implemented is to set the beam splitter at a dark fringe of the light path, setting $\phi_0 = 0$, so that no light passes if no signal is present, making the interferometer a *null instrument*. A few more precautions must be taken to actually detect the signals, since at this working point

$$\left(\frac{\partial P}{\partial \phi}\right)_{\phi_0=0} = 0 \tag{2.4}$$

Since $\Delta\phi_{gw} = \mathcal{O}(h)$, the change at the power output induced by a gravitational wave will be $\Delta P = \mathcal{O}(h^2)$ at the dark fringe, which is basically invisible. The strategy used to induce a linear change in power output at the dark fringe relies on introducing sidebands in the laser light [1]. The electric field of a laser on which a phase modulator operating at frequency Ω_{mod} is applied will have the form

$$E_{in} = E_0 e^{-i\omega_L t + \Gamma \sin(\Omega_{mod} t)} \tag{2.5}$$

Where Γ is the modulation depth. The two sidebands will be at $\lambda_\pm = \lambda_L \pm \lambda_{mod}$ By taking $\Gamma \ll 1$ One can expand this in linear order, and the power output can be dived into a therm due to the carrier (E_c) and two due to the modulation ($E_\pm$). These will be

$$\begin{aligned} E_c &= -iE_0 J_0(\Gamma)\sin\left[2\pi\frac{L_x - L_y}{\lambda_L}\right] \\ E_\pm &= \mp iE_0 J_1(\Gamma)\sin\left[2\pi\left(\frac{\Delta L}{\lambda_L} \pm \frac{\Delta L}{\lambda_{mod}}\right)\right] \end{aligned} \tag{2.6}$$

If $L_x = L_y$ then both the carriers and the sidebands will be on the dark fringe, but whenever $\Delta L = n\lambda_L$ the carrier will be on the dark fringe while the sidebands will not. This effect is called the Schnupp asymmetry.

If a + polarized gravitational wave perpendicular to the interferometer is passing, $\Delta L_{gw} = \Delta L + Lh(t)$. By defining $E_s = E_+ + E_-$ the total power output will be $\Delta\ E_{tot}^2 = (E_c + E_s)^2$ of an interferometer working in the Schnupp asymmetry condition. There will be three terms in the equation: a E_c^2 therm that contains a contribution from the gravitational wave of order $\mathcal{O}(h^2)$, which is too small to be observed; a E_s^2 term that will be of order $\mathcal{O}(1) + \mathcal{O}(h)$ that overwhelms the signal; finally, a therm $E_c E_s$ oscillating at Ω_{mod} frequency, originating from the beatings between the carrier and the sidebands, that will be of order $\mathcal{O}(h)$. So, by applying a demodulator at the output and retrieving the therm at frequency Ω_{mod}, the detector can achieve the low noise condition of the dark fringe without having to sacrifice the linear order in $h(t)$.

To optimally design an interferometer to observe gravitational waves, one could find the optimal length for the Michelson's arms by looking at equation 2.2. Here the therm $L\mathrm{sinc}(\omega_{gw}L/c)$ needs to be maximized. So the optimal arm length will depend on the frequency being targeted, and this relation will be $\omega_{gw}L/c = \pi/2$ or $L = \lambda_{gw}/4$ [2]. When this condition applies, the gravitational wave will act on spacetime with the same sign during the whole time of flight of the photon, maximizing the phase displacement. By designing an interferometer set to maximize the effect of gravitational waves at $f_{gw} = 100$Hz one finds that the arms should have length $L = 750km$. This kind of length is impractical for ground-based interferometers, but modern detectors can work around this issue by virtually elongating the arms with optical cavities. The working principle of the Fabry-Perot cavities that are used in Advanced Virgo and every other modern gravitational wave interferometer is presented in the next section.

2.1.2 Fabry-Perot cavities

Fabry-Perot (FP) cavities are implemented in Michelson interferometers by placing input mirrors between the beam splitter and the arms. If the input and output mirrors have high reflectivities, the photons will bounce back and forth along the arms. The effect of gravitational waves on $\Delta\phi_{mich}$ will therefore be additive for each round trip, maximizing the effective length of the arms.

Let's take a look at the problem more quantitatively, by taking as an example a single FP cavity. Defining r_1, t_1 and r_2, t_2 as the reflectivity and transmissivity of the input and the end

mirrors, if a beam with power P_0 is shone in the cavity, the power stored inside will be [1]:

$$P_{cav} = \frac{t_1^2}{1 - (r_1 r_2)^2 - 2 r_1 r_2 \cos(2kL)} P_0 \tag{2.7}$$

When $L = n\pi/k = n\lambda/2$ the cavity is at resonance and the power contained in it is at its maximum. This happens when the entering electromagnetic field is in phase with the one inside the cavity. If the length of the cavity is perturbed away from resonance, a drop in the power inside (and subsequently the power output) will be observed. The sharpness of this resonance peak with respect to the cavity length is determined by the finesse parameter $\mathcal{F}$, which, with the assumption of really small cavity losses, is the ratio between the distance of resonant peaks and the peak width at half maximum power:

$$\mathcal{F} = \frac{\pi\sqrt{r_1 r_2}}{1 - r_1 r_2} \sim \frac{2\pi}{losses} \tag{2.8}$$

From the finesse, it is possible also to obtain the storage time of a cavity, or in other words, the mean time a photon spends inside it [2]:

$$\tau_s \simeq \frac{L\mathcal{F}}{c\pi} \tag{2.9}$$

For Advanced Virgo, the finesse of the arms is set at roughly 440 [21], meaning that the storage time is around $\sim$ 2ms. This is roughly a quarter of a period of a 100 Hz oscillation, which as discussed in the previous section is the optimal time of flight of a photon for observing gravitational waves at 100 Hz frequency in a Michelson Interferometer.

If the length of a cavity at resonance is perturbed with a change ϵ, it can be demonstrated [1] that the derivative of the phase of the output beam with respect to the ϵ will be

$$\frac{\partial\phi}{\partial\epsilon} \simeq \frac{2\mathcal{F}}{\pi} \tag{2.10}$$

which means that small changes in cavity length can get greatly amplified in the output. If the arms of a regular Michelson interferometer are replaced by Fabry-Perot cavities, then these can radically improve the sensitivity of the instrument, basically multiplying the phase change induced by gravitational waves on the arms by the finesse parameter. This will have the effect of simulating the sensitivity that a much larger simple Michelson interferometer would have. At first approximation, one could think that by setting the finesse to an arbitrarily large value, one could simulate an arbitrarily large interferometer, but the higher power stored in the cavities will put higher and higher radiation pressure on the mirrors. This introduces noise in the measurement, limiting the maximum achievable finesse. This radiation pressure effect will be explored in sec. 2.4.

The capabilities offered by Michelson interferometers equipped with Fabry-Perot cavities for observing such small changes is the main reason why they are the instrument of choice for the measurement of gravitational waves.

2.2 Advanced Virgo

This section will show how all of the ideas introduced in the previous parts of the chapter are implemented in the actual Advanced Virgo interferometer. A schematic of the instrument is shown in fig. 2.2. Located in the countryside of the Italian town of Cascina, near Pisa, the instrument is a Michelson interferometer with 3 km long arms, both equipped with Fabry-Perot cavities [22]. The instrument is comprised of four main buildings: The Central Building (CEB), the Mode Cleaner Building (MCB), the North End Building (NEB) and the West End Building (WEB) The rest of the section will describe the most important of the optical components of the interferometer.

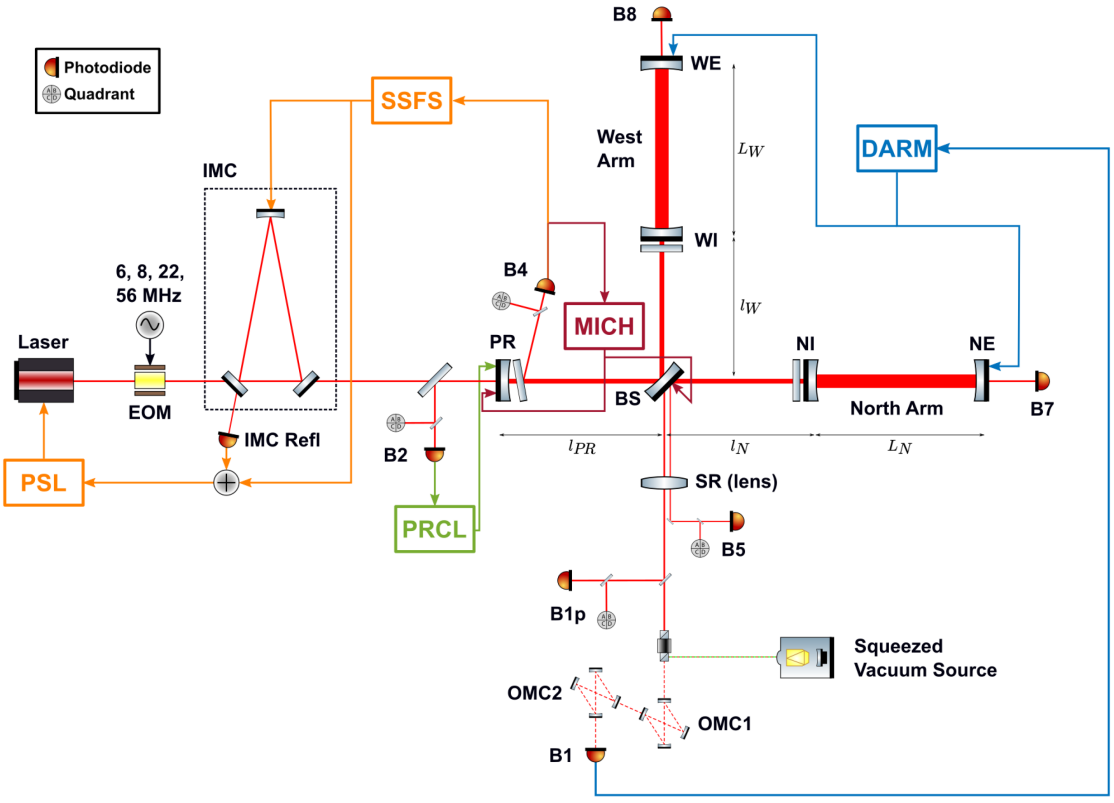

Figure 2.2: An overview of the Advanced Virgo optical layout during O3 [21]

- **Laser:** The laser used in Advanced Virgo is a 25 W PSL infrared laser (1.064μm) [21]. The beam is sent to an Electro-Optical Modulator (EOM) mounted on the optical Input Bench

(IB) that introduces the sideband needed to make the detection possible when the beam splitter is at the dark fringe.

- **Input mode cleaner (IMC)** A 144 m cavity formed by two mirrors mounted on the IB and a mirror in the Mode Cleaner (MC) building. This cavity is resonant only to Gaussian-shaped beams so that every other mode will be suppressed in the output. The clean Gaussian laser is needed since the FP cavities in the rest of the interferometer are tuned to work with Gaussian beams only: the end mirrors in the arms are spherically shaped to create a surface of constant phase for the incoming EM field. With a normal flat mirror, the beam would get wider and wider after each bounce, but this special surface will reflect the Gaussian beam back into its shape, avoiding lateral dispersion. If modes other than the Gaussian enter the interferometer, the dispersed beams could get randomly scattered around the arms, causing unwanted noise if they reach the photodetector.

- **Power Recycling mirror (PR)** This mirror, placed between the IB and the Beam splitter increases the effective power circulating inside the rest of the interferometer. In normal operating conditions, when no signal is present and all of the cavities are in resonance, when the BS is at the dark fringe with respect to the output port, all of the light coming from the arms will get reflected back towards the input. From the point of view of the IB, the rest of the interferometer is acting just as one highly reflective mirror. By introducing another highly reflective mirror at this point, the PR mirror, a new FP cavity can be formed that will make the power in the whole interferometer grow considerably. When in resonance, this cavity increases the effective power of the laser to 1 kW. [21].

- **North and West cavities:** These are the most important cavities of the interferometer since their difference will be the only one sensitive to the passage of gravitational waves. They are each 3 km long and are oriented roughly in the North-South and East-West directions. Their input mirrors are called the North Input (NI) and West Input (WI) and are situated inside the CEB, the end mirrors are referred to as the North End (NE) and West End (WE) and are situated at the NEB and WEB building respectively. When in resonance, the power inside these cavities can reach 130kW [21].

- **Signal Recycling mirror (SR)** set between the BS and the photo-detector, It will be part of a cavity for O4 that will improve detector sensitivity. For a stand-in, during O3 the first lens of the output bench was put in its place.

- **Output Mode Cleaner (OMC)** two small cavities (OMC1 and OMC2) located on the optical Output Bench (OB) that have a job similar to the IMC, this time to avoid spurious modes that can be generated by inaccuracies in the alignment of the mirrors reaching the detector output.

- **Squeezing bench**: An optical bench that introduces a squeezed vacuum state in the output beam [23]. This is put in place to reduce the quantum shot noise that is introduced by the random fluctuations of the local electromagnetic field at the output port of the interferometer, which will be discussed in more detail in sec. 2.4. Usually, these fluctuations are symmetric in the amplitude and phase quadratures of the EM field, but one can reduce the noise in one quadrature while amplifying the noise in the other, always following the Heisenberg Uncertainty Principle. By reducing the uncertainty in the phase while amplifying the one in the Amplitude, the sensitivity of the detector is improved at high frequencies while being worsened at lower ones, but since low-frequency noise is well dominated by other sources, this is a worthwhile trade-off. For O4, frequency-dependent squeezing will be implemented [18].

2.3 Superattenuators

Superattenuators are the current mirror seismic isolation system in the Advanced Virgo Interferometer [24]. These are complex machines fundamental for the correct operations of the instrument, acting both as seismic dampeners and as tools to move the mirror to its working position. If transferred directly to the mirrors, the normal ground movement at the site could completely overwhelm any gravitational wave signal, being several orders of magnitude louder. So the Superattenuator operates both passively and actively to reduce the seismic noise by more than 10 orders of magnitude in the sensitive frequency range of the detector. The rest of this section will give a brief overview of the physical principles and the design choices of the Superattenuators.

The main attenuation component of the Superattenuator is the pendulum chain. When considering mechanical excitation that acts on the suspension point of a pendulum as an input(x_{in}), and the suspended mass movements as an output (x_{out}), the whole pendulum can be considered a mechanical second order low pass filter. A pendulum that resonates at ω_0 will have a transfer function in the horizontal plane for frequencies $\omega \gg \omega_0$:

$$\frac{|\tilde{x}_{out}(\omega)|}{|\tilde{x}_{in}(\omega)|} \simeq \frac{1}{\omega^2} \tag{2.11}$$

Meaning that frequencies much higher than ω_0 will be greatly reduced in intensity before getting to the mass. The cutoff frequency is determined by the pendulum length, so by tuning this parameter one could obtain great dampening for some desired frequencies. Moreover, if a second pendulum is suspended from the hanging mass of the first, the dampening effect will be compounded, creating an effective fourth-order low pass filet (dampening like $\sim 1/\omega^4$), and this process can continue by hanging more and more pendulums. This is the reason why pendulums

were chosen to be the mean through which passive seismic dampening is achieved in Advanced Virgo. Since in the Superattenuators the wire length is 1.2 m, the resonance will be at 0.5Hz, which is well below the lowest detectable frequency of the detector. There exist two types of Superattenuators chains in Virgo, one with 6 filters and 10 meters in height, reserved for the most important mirrors: NE, NI, WE, WI, BS, PR and SR, while shorter 3 filters, 4.5 meters towers are installed for the IB, MC, and OB Superattenuators [25]. In fig. 2.4 a schematics of the long chain Superattenuator is present. From now on the description of the Superattenuators will focus only on the long-chain configuration.

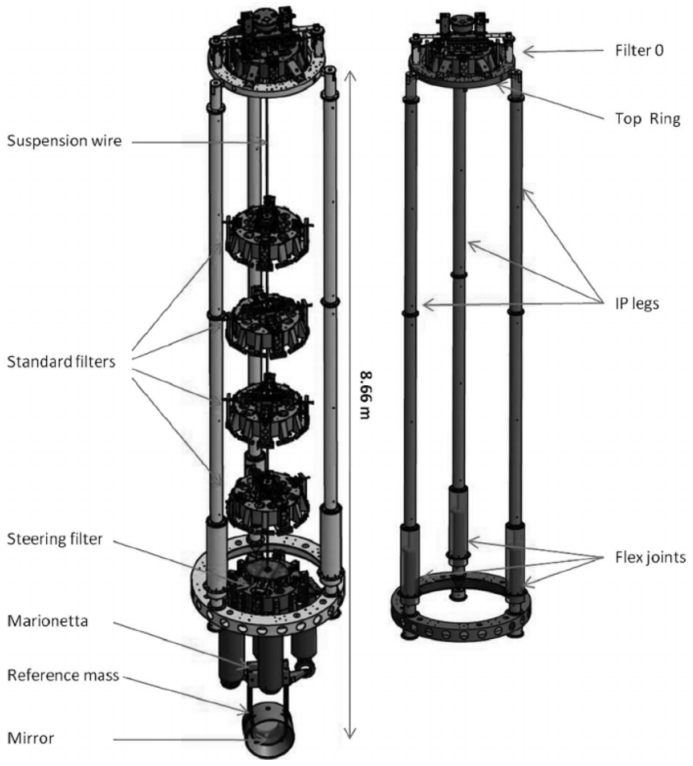

Figure 2.3: Advanced Virgo Superattenuators schematics for the long chain configuration. On the right, the whole Superattenuator is pictured while on the left only the pre-isolation stage is present.[26]

The filter chain is linked to the ground by a pre-isolation stage, where the main body is a three-legged structure 8 meters in height, called the Inverted Pendulum (IP) [24]. IPs are rigid-body pendulums where the mass is at the unstable equilibrium point. Like their non-inverted counterparts, these can act as low-pass filters. In the small oscillations regime, an IP with length L that is supporting a mass m and that has a flexible joint with stiffness k will have a resonant frequency

$$f_0 = \sqrt{\frac{k}{mL^2} - \frac{g}{L}} \tag{2.12}$$

The fact that f_0 is inversely proportional to the square root of the suspended mass, means that IPs can achieve normal modes with really low resonant frequencies. For the Superattenuators, since they present a total mass of $\sim 10^3$ kg, the resonance frequency is around 30 mHz. So by attaching the pendulum chain to the top of the IP, great amounts of dampening can be achieved even before the main seismic suppression instrument. The IP also provides a soft suspension point to the mirror, making the forces required to control the position much weaker. The IPs legs are joined together both at the top and at the bottom by metal rings 1.5 m in diameter, called the Top Ring and the Base Ring. The Base Ring is connected to the ground by three vertical actuators, that provide both active vertical attenuation and the control needed to keep the IP near its unstable equilibrium point. The top ring acts as a platform from which the first stage of the pendulum chain can be attached, the so-called Filter0 (F0) which hosts a variety of instruments. Positioned both vertically and horizontally in a pin-wheel configuration, Linear Variable Differential Transformers (LVDT) and accelerometers are the sensors that inform the system on the F0 position [25]. Even if these are not oriented as the normal Cartesian coordinates, virtual sensors are created by having a linear combination of their outputs, to have measurements in the x, y, z and θy d.o.f., where z is the local direction of the laser, x is the other horizontal component, y is vertical and θ_y is the rotation around the vertical axis. The outputs of the sensors are used in the active dampening strategy, where coil-magnet pairs will act on the suspension point to suppress the low-frequency oscillations that the pre-isolation stage let trough, so under and around its ω_0, and also they will act to suppress the normal modes of the SA.

Each filter of the chain is a drum-shaped contraption that weighs around 120 kg, hung from the previous ones by a steel wire. At the connection point, a set of triangular blades placed on the base of the filter provide the vertical attenuation [27]. The blades act as springs that filter out the residual motion in the vertical direction. The number of blades depends on the load that the wire needs to withstand (12 at the F0, just 4 at the last stage), so that the resonant frequency of the spring, which depends on the load attached to it, is kept the same all along the chain. Magnetic anti-springs are also used in the vertical attenuation to bring down the resonant frequency from 1.4Hz typical of the blades down to 0.5 Hz so that the attenuation will be isotropic in the 3 DoFs. Achieving vertical attenuation is quite important since the vertical component can couple to the horizontal one measured by the laser: The curvature of the earth changes the local vertical direction, and over the 3 km of the Virgo arms this effects is quite noticeable: if the mirrors experience a vertical displacement δ_y the corresponding δ_z in the laser direction will be $\delta_z \sim \delta_y 4.3 \times 10^{-4}$.

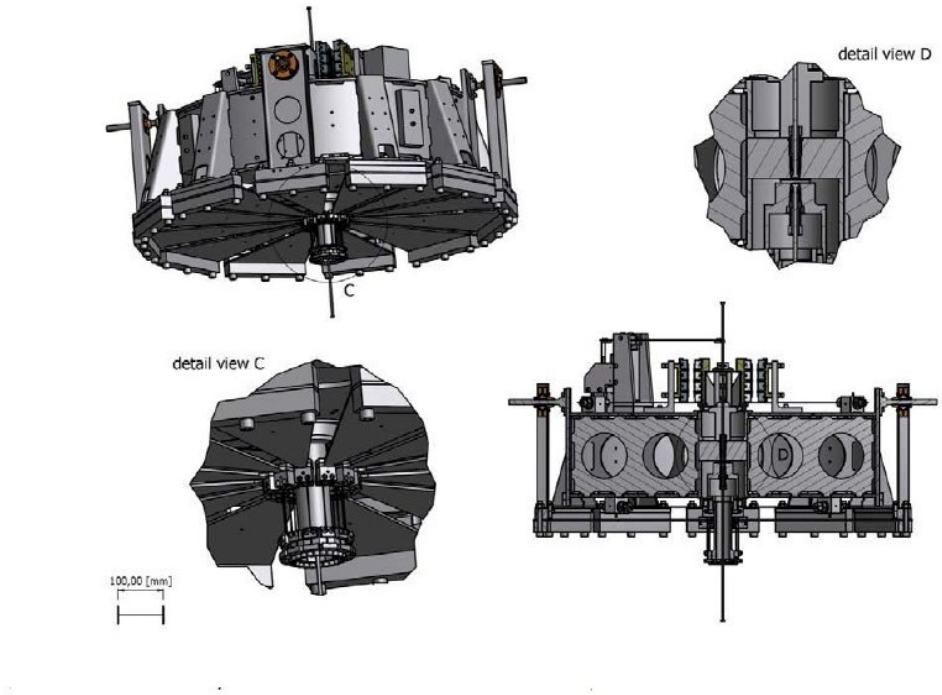

Figure 2.4: Schematic of the standard filters of the Superattenuaotrs in Advanced Virgo [25]

The last filter is called for historical reasons Filter7, where sensors and actuators are used to control the angular DoFs ($\theta_x, \theta_y, \theta_z$). A metallic structure, called the Marionetta, is hung from the F7 and provides an anchor point for the 4 fused silica fibers from which the mirror is hung. It also acts as a platform where to put the magnetic coils that act directly on the mirror to achieve the finest positioning in the whole system.

A lot of work has been done to characterize the transfer function of the Superattenuator as a whole, with great success [25], [28] but due to their complexity, sometimes simulations can quite differ from reality, especially in the most extreme situations. Non-linear dependencies and the really small nature of the perturbations make this task quite a difficult one. But the vast amounts of data acquired during the operating years of the interferometer can be looked at to study the response of the Superattenuator to seismic excitation, and Machine Learning can learn the non-linearities of the system. This is the reason why this framework was chosen to explore this problem.

2.4 Noise sources in Advanced Virgo

Like every other high-precision instrument, noise has a great impact on the final measurement in Advanced Virgo. To get useful science data out of the detector, the noise has to be accurately

studied and characterized. The interferometer has a complex noise landscape, happening in transient, stationary and quasi-stationary regimes. This section will be focused on giving a brief description of the stationary noise that affects the detector. For Advanced Virgo the noise has a strong frequency dependence since in different parts of the spectrum different sources dominate the measurement uncertainty. This can be explored by calculating the Power Spectral Density (PSD) of $h(t)$, which is a measure of the power contained in the signal for each frequency band. The power spectral density for a stationary process $x(t)$ is defined by:

$$S_x(f) = \lim_{T\to\infty} \frac{2}{T} \left| \int_{-T/2}^{T/2} x(t) e^{-2\pi i f t} dt \right|^2 . \tag{2.13}$$

By integrating the PSD in frequency one obtains the power contained in that frequency band only. In fig. 2.5 an approximation of Advanced Virgo PSD is present, where the contributions of each of the sources are highlighted, and in fig. 2.6 the real PSDs of the LIGO-Virgo collaboration detectors are plotted. A brief description of the main noise sources in the Advanced Virgo interferometer will now be given.

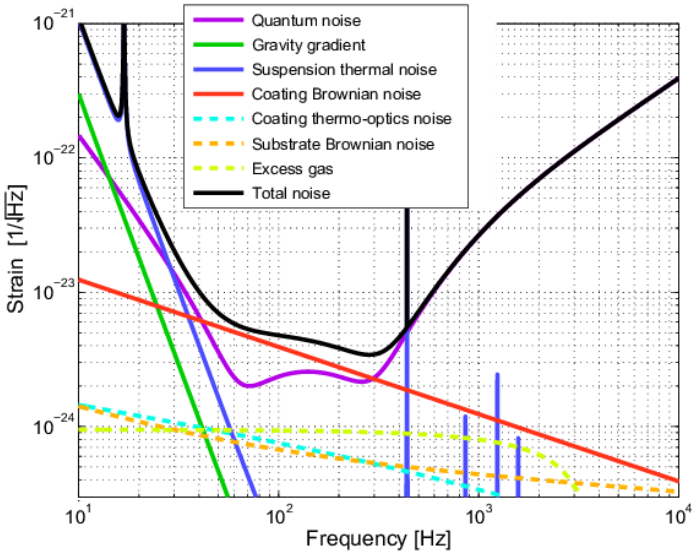

Figure 2.5: Simplified PSD of the Advanced Virgo Noise budget, with some of the known causes plotted on the graph[29]

2.4.1 Quantum noise

Quantum noise is at the moment the most predominant noise source that dominates the high frequencies. It originates as a consequence of the Heisenberg uncertainty principle that induces random quantum fluctuations in the electromagnetic field near the interferometer output port

[2]. This noise is divided into two main components:

- **Shot noise**: It originates from the probabilistic distribution of the number of photons that reach the photodiode, resulting in a measurable power fluctuation. The photon count N follows a Poisson distribution with $\sigma = \sqrt{N}$ and so the power fluctuation observed over a period T will be

$$\Delta P = \sqrt{P_0 \hbar \omega_l / T} \tag{2.14}$$

  At the photodetector, the effect of a gravitational wave will induce power fluctuations linear to P_0, so by considering the signal-to-noise ratio, the PSD of the shot noise will be

$$S^{1/2}(f) \propto \frac{1}{\mathcal{F} L \sqrt{P_0}} \sqrt{1 + (f/fp)^2} \tag{2.15}$$

  Where f_p is the cutoff frequency of the Fabry-Perot cavity. This noise increases with higher frequencies, but it can be lowered by increasing the power stored in the cavities, by either increasing P_0 or the finesse $\mathcal{F}$.

- **Radiation Pressure** Photons hitting the mirrors exert on them significant radiation pressure. Due to the probabilistic distribution of photons discussed in the previous section, this force is not constant and induces random oscillations on the mirror. This force is, therefore, proportional to P_0 and, after being filtered through the transfer function of a free-falling mass M, the PSD equation of the noise at the detector will be

$$S^{1/2} \propto \frac{\mathcal{F}}{M L f^2} \sqrt{P_0} \sqrt{\frac{1}{1 + (f/f_p)^2}} \tag{2.16}$$

  Contrary to shot noise, radiation pressure noise is louder at lower frequencies. To decrease these disturbances, the strategies implemented in modern interferometers is having mirrors with higher mass [18] [16]. Contrary to the previous case, an increase in the stored cavity power increases the loudness of the noise, so to get the best performance out of the interferometer, a compromise must be made. At the current time, Shot noise is the most limiting factor, since at lower frequencies seismic noise still dominates over radiation pressure. A way to limit both kinds of noises is the implementation of frequency-dependent light squeezing and will be installed during the O4 run [23].

2.4.2 Thermal noise

The mirrors and the Superattenuator towers are kept at room temperature. This means that Brownian motion can induce random fluctuations in the mirror position. The power spectrum of the displacement due to Brownian motion is linked to the frequency like $\tilde{x}(f) \propto \sqrt{T}/f$, meaning that it will impact the lower frequencies more and it increases with temperature. This noise

comes from 2 main contributions, the suspension's thermal noise, and the mirror's thermal noise.

The Superattenuator pendulum chain limits the thermal noise from every filter but the last, which gets transferred directly to the mirror. Thermal vibrations that excite the resonant modes of the wires that hold the mirrors in place can also occur, and cause high-frequency lines starting at 300Hz and going up to a few kHz, called violin modes. [30]

The mirror's thermal noise comes from different contributions since it can either come from the Brownian motion of the reflective coating of the mirror or the bulk when oscillations due to thermal expansion or changes in the refractive index of the mirror can be seen. These fluctuations can be avoided by taking the mirrors to cryogenic temperatures, and this is the focus of new-generation interferometers like the recently operational KAGRA observatory and the future Einstein Telescope [31].

2.4.3 Seismic Noise

Seismic noise is the main limiting factor at low frequencies for current ground-based interferometers. The Seismic motion spectrum usually takes the form of:

$$S^{1/2}(f) \simeq \frac{\alpha}{f^2} \tag{2.17}$$

Where α is a quantity typical of the site and season, and at Virgo is usually around $\alpha \simeq 10^{-7} m\text{Hz}^{3/2}$ [25]. This noise is more than 10 orders of magnitude greater than the displacement caused by gravitational waves. This noise is not only capable of completely masking these weak signals at low frequencies but can also make the mirror drift away from its operating position at the resonance of the FP cavities, incapacitating the measurement capability of the interferometer as a whole. The seismic landscape at Virgo will be discussed in more detail in sec. 4.2.

2.4.4 Newtonian noise

Seismic noise can still creep into the measurements, bypassing the suspensions, in the form of Newtonian noise. Oscillations in the local distribution of mass around the mirrors can modify the gravity gradient, inducing unwanted movements. Newtonian noise is not possible to mitigate, since any suppression strategy would mask gravitational waves as well. But, since these gravity gradients are a strictly local phenomenon, Newtonian noise can be avoided by moving the interferometer to a place where the gravity gradient is smoother and less prone to oscillations, like underground or in outer space. These strategies will be implemented by future detectors like ET, an underground gravitational waves interferometer planned by the European Union [31], or the space mission LISA [14].

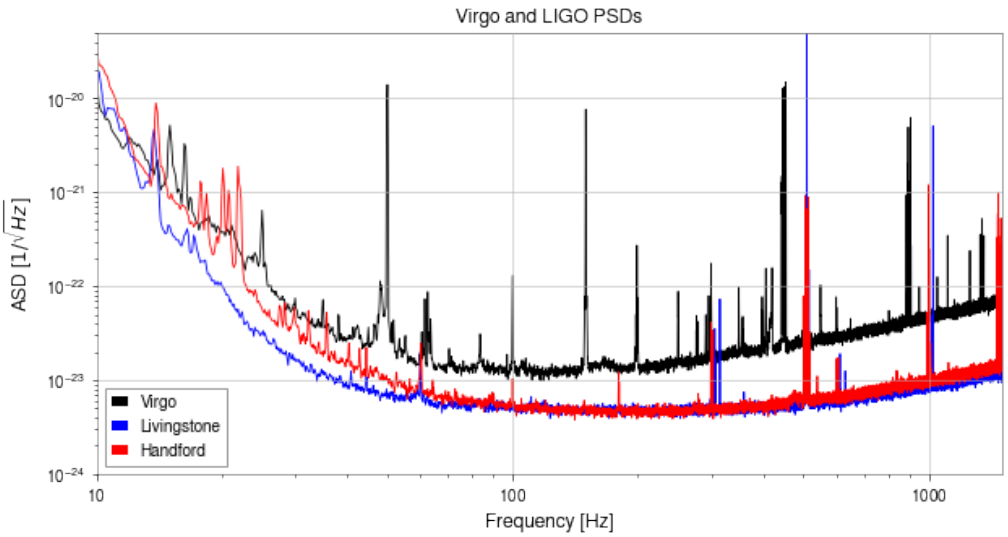

Figure 2.6: Real LIGO and Virgo Power spectral densities during O3b on 06 January 2020, integrated over 500 seconds. All of the interferometers share similar limitations so their PSDs have roughly the same shape

2.5 Transient Noise (Glitches)

While noise sources listed so far are continuous in nature and limit the overall sensitivity of the instruments, interferometers are also impacted by transient and short-duration sources of noise, called glitches. These appear as high Signal to Noise ratio events that can obscure or even mimic the passage of gravitational waves. The understanding of this kind of noise is fundamental to exclude them in the search for gravitational signals, and lowering the false alarm rate that can hinder the low latency pipeline needed for multi-messenger astronomy.

A lot of work has been done to study and mitigate glitches, but most of them are not well understood yet, and many of the causes remain unknown. A lot of distinct classes of glitches have been found and they seem to have different origins; citizen science initiatives like Gravity spy [32] or Gwitch hunters [33] have found 20 different classes of glitches, that range in the whole frequency spectrum. Some correlations have been found between glitches and auxiliary channels, which helps the vetos algorithm to exclude them faster from gravitational wave candidates. Glitches concerning the lower part of the spectrum, for example, classified as low-frequency lines, low-frequency bursts and scattered light, seem to be the ones more linked to unwanted oscillations of the instrument's components. Scattered light glitches, in particular, have a peculiar "banana" shape and they seem to be mostly linked to rouge light beams outside of the main FP cavities that hist some components of the interferometer and get backscattered into

the main beam, and since non-isolated interferometer components are more likely to be coupled to seismic noise, these disturbances can appear in the main channel [34]. Low-frequency glitches can be caused by bad weather, that makes the buildings shake, or by local activities causing unpredicted ground motion, but a lot of other causes may be the culprit. These kinds of glitches will be of interest to study since they can be linked to the seismic activity at the site, the study objective of this thesis.

2.6 Control of the Interferometer

To get the interferometer to its working condition, a lot of steps must be taken. The most important task needed to get the detector ready is to tune all of the optical cavities to resonance, and subsequently keep them stable. The process of finding and acquiring resonance for a cavity is called locking. A complex control feedback loop, based on sidebands injected in the laser light, measures the error between the current and the wanted position of the mirror. Subsequently, these error signals instruct piezometric actuators that act on the 6 degrees of freedom of the mirrors to move them to their working point. The steps needed to get the interferometer in the so-called "Science mode" are described by the ITF_index auxiliary channel, which indicates how many steps down the checklist have been taken to get the detector to the working point. The interferometer during O3 was at its full capabilities when the index was at 170 and the process to get all of the steps of the checklist complete can take a few hours. The critical steps are organized in the flow chart in fig. 2.7 and are described in [21].

The process starts by locking the North and West arm cavities, first independently and then syncing them together to have the correct working point at the beam-splitter. After the mode cleaner cavity is locked, the Second Stage of Frequency Stabilization (SSFS) can be turned on, which consists in using the Common Arm Length (CARM) measurement, which is the mean length of the West and North cavities, to have fine control over the laser frequency instabilities, since this DoF is really sensible to them, to stabilize the laser output. The Power recycling mirror can now be brought to its operating point, so now the power in the whole Interferometer increases dramatically. After having reached this point, the mirrors do not need to make dramatic movements anymore, so the actuators are switched to low noise mode 1, which limits their dynamic range and subsequently their electrical noise. The output mode cleaner can now be locked and at this point, the interferometer is finally able to measure the Differential Arm length (DARM) Dof, which is the one sensible to gravitational waves. To further reduce the noise the dynamic range of the actuators is lowered again, to low noise mode 2 and then 3 and finally, the fully operational science mode condition can be achieved by injecting the squeezed vacuum in the output port of the Interferometer.

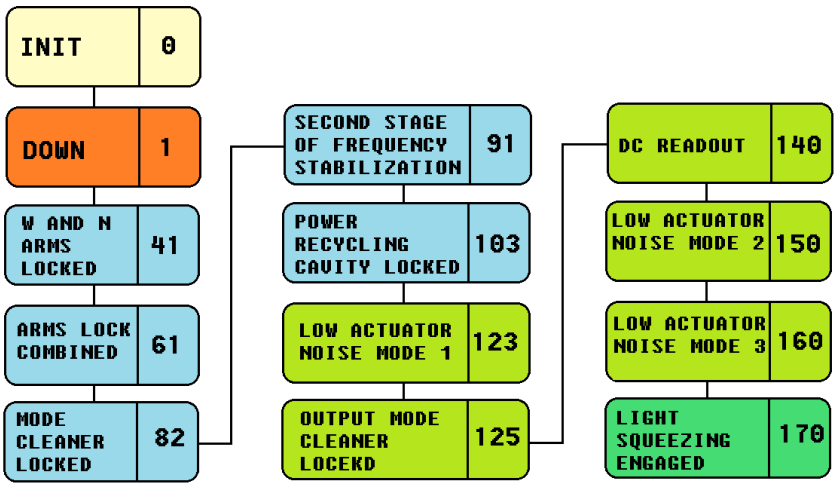

Figure 2.7: An overview of Advanced Virgo steps to get to science mode.

Since the actuators of the Superattenuators change mode of operation based on the interferometer status, this will be a useful quantity to check during the analysis, and also this will be an indication of the quality of detection in general during these events.

Summary

In this chapter, the Michelson interferometer was introduced, which is the instrument of choice when trying to detect gravitational waves. It was then discussed how a simple Michelson could be improved upon, for the specific task of detecting gravitational waves, by implementing Fabry-Perot cavities. How all of these concepts are actually implemented was given in the next section, which described the actual layout of the Advanced Virgo interferometer, with a focus on the seismic suppression system at Virgo, the Superattenuator. An overview of the main noise sources that limit the sensitivity of the interferometer was given next and finally, the chapter is concluded by describing the procedure through which the interferometer is brought to the data acquisition phase.

Chapter 3

Machine Learning for gravitational waves physics

In recent years, Machine Learning algorithms have become more and more prominent in scientific research [35]. In particular, their ability to learn arbitrarily complex tasks from vast amounts of data is the main reason why the project hinges on these techniques to find patterns through the complexity of Virgo auxiliary channel data. This chapter will be an overview of the algorithms and techniques used in the project. Sec. 3.1 will serve as a broad overview of the concepts behind machine learning and its principles. A summary of how machine learning has already been successfully implemented in different fields of gravitational wave astronomy is given in sec. 3.2. In sec. 3.3 a description of the dimensionality reduction algorithm called Variational Autoencoder is given since it will be used quite extensively in the final project and then the clustering algorithm Gaussian Mixture Model will be described in sec. 3.4.

3.1 Introduciton to Machine Learning

Machine Learning refers to the broad field of algorithms that are capable of learning to perform tasks just from the observations of large amounts of data. These tasks can be divided into the three broad categories[36]:

- Supervised Learning, where an algorithm will try to approximate a complex unknown function $f(\mathbf{x}) = \mathbf{y}$ with another function $\tilde{f}(\mathbf{x})_\theta = \tilde{\mathbf{y}}$, that depends on the θ parameters, by training on a sample dataset of input-outputs pairs $\mathbf{X}$ and $\mathbf{Y}$. This category includes classification and regression tasks.

- Unsupervised Learning, where the algorithm will try to learn a generative model for the dataset $\mathbf{X}$. The algorithm will try to recreate samples similar to $\mathbf{X}$, where each one is generated starting from a latent variable vector $\mathbf{z}$. The learned function $\tilde{f}(\mathbf{z})_\theta = \tilde{\mathbf{x}}$ is

considered to be an approximation of the real unknown function that generated the data. This category comprises tasks like clustering and the generation of synthetic data.

- Reinforcement Learning, where the algorithm will try to learn a policy function to decide actions to take to achieve a certain objective based on observations of the environment.

This thesis will mostly focus on the unsupervised learning framework since these methods are considered the most flexible and do not rely on the presence of labeled data. Their main capabilities rely on finding statistical properties inherent in the dataset, which is the final objective. For the rest of the section, the focus will be on just unsupervised learning.

To find the parameters θ that define these generative functions, optimization algorithms are run to minimize some penalty $\mathcal{L}(\mathbf{y}, f_\theta(\mathbf{x}))$, called the loss function, which measures the distance between the algorithm output $f_\theta(\mathbf{x})$ and the desired outcome $\mathbf{y}$. For generative models, this penalty is usually a distance measure between the original dataset and the one generated by the algorithm. For example, the Mean Squared Error (MSE) between the input and the outputs can be used. This is just the mean L^2 norm:

$$\mathcal{L}(\mathbf{y}, f_\theta(\mathbf{x})) = \frac{1}{N}\sum_{i=1}^{N} \|\mathbf{y_i} - f_\theta(\mathbf{x_i})\|^2. \tag{3.1}$$

Another common loss, which is mostly used in classification tasks but has also seen successful uses in VAEs is the Binary Cross Entropy (BCE):

$$\mathcal{L}(\mathbf{y}, \tilde{\mathbf{y}}) = \sum_{i=0}^{N} -y_i \log(\tilde{y}_i) - (1 - y_i)\log(1 - \tilde{y}_i) \tag{3.2}$$

where $\tilde{y}_i$ is the prediction of the network. BCE only works if y and $\tilde{y}$ are bounded between 0 and 1 since they should represent probabilities. In generative models, one could bound the variables between 0 and 1 and consider the output as the probability that the j-th pixel is "on". Both MSE and BCE were tried in this thesis.

The optimization of the loss function is achieved in practice through iterative methods based on the minimization of the gradient of $\mathcal{L}$, called gradient descent methods. Starting from a random initialization of the parameters $\theta^{(0)}$, the output is calculated $f(\mathbf{X})^{(0)}_\theta = \tilde{\mathbf{Y}}^{(0)}$. Standard gradient descents methods iterate until convergence to find a local minimum of the function, updating the parameters θ as

$$\theta^{(t+1)} = \theta^{(t)} - \eta \triangledown \mathcal{L}(\mathbf{Y}, f(\mathbf{X}; \theta^{(t)})) \tag{3.3}$$

Where η is the learning rate parameter, that determines how big of a leap through parameter space the algorithm will take at each iteration. The choice of this parameter is critical since a

large learning rate will lead to an unstable algorithm, that will jump over local minima, while a small learning rate will lead to a more stable algorithm, but that will take a long time to converge. When properly tuned and applied to convex continuous functions, this method will eventually reach convergence on a local minimum [36]. "Vanilla" implementations of gradient descent are rarely used in practice, since the optimization might take a long time and might also get stuck in a local minimum far from the global one. Many variations of this algorithm exist to overcome these issues. One of these is the Stochastic Gradient Descent (SGD): instead of computing the Gradient over the whole dataset, the algorithm computes it over small randomized batches of samples and runs the optimization for each batch. By using this technique, some stochasticity given by the random nature of the samples is introduced in the algorithm. This helps the optimizer to avoid getting stuck in a local minimum. Another common implementation is to add a "momentum" term to the weights update. The momentum is calculated by taking a running average of the previous computations of the gradient, which is then added with a weight parameter $0 < \beta < 1$ to the update parameter step. The equation for the optimization will now be

$$\begin{aligned} v^{(t)} &= \beta v^{(t-1)} + \nabla \mathcal{L}(\mathbf{Y}, f(\mathbf{X}; \theta^{(t)})) \\ \theta^{(t+1)} &= \theta^{(t)} - \eta v^{(t)} \end{aligned} \tag{3.4}$$

This is again done to avoid the algorithm getting stuck in a local minimum, and also helps with the "vanishing gradient" problem. This arises whenever the algorithm encounters parts of the parameter space where the loss function does not show any improvements, resulting in gradients with really low values. But, if the gradient inclination is consistent, then the momentum will compound over each iteration, making the learning jump longer and longer, adapting to the current landscape. Successful implementations of momentum in gradient descent algorithm are RMSprop [37] and Adam [38].

Since these optimization methods rely on the computation of the gradient, a function $f_\theta(\mathbf{x})$ that allows a fast and stable calculation of this quantity is required. At the same time, f_θ needs to be flexible and expressive enough to approximate the unknown real output function. Biology-inspired algorithms called neural networks can accomplish both of these tasks and are used extensively in many machine learning applications. In these algorithms, the input vector $\mathbf{x} \in R^n$ is multiplied by the weight matrix $W^1_{n\times k}$, and summed to a bias vector $b^1 \in R^k$. The now transformed vector is referred to as the hidden layer, and the dimension k is user-defined, the bigger the k, the more expressive the network will be. A non-linear function σ is then applied to the hidden layer. This is referred to as the activation function. This step is repeated as many times as the user desires until the last step, which produces an output vector $\tilde{\mathbf{y}} \in R^m$, on which the loss function will be calculated. The optimization algorithm will then tune the weight matrices W^1,W^2 ... and the biases vectors b^1, b^2... to accomplish the minimization of the loss. It can be proven that a neural network with just one hidden layer with an arbitrary k and a

non-linear activation σ can approximate any function $f(\mathbf{x}) : R^n \times R^m$ [39], so this makes neural network incredibly flexible architectures, which was the first requirement. Since the gradient depends on the activation function σ, it can be chosen so that this calculation is as simple as possible. The most popular choices for this function in the literature are Sigmoids, arctangent and the widely used ReLU function $\max(0, x)$ where the computation of the gradient is the fastest, being 1 if the input is positive and 0 otherwise. One other technique that makes neural networks the architecture of choice for many machine learning applications is how relatively easy is to calculate the contribution of the gradient of every θ_i parameter, even for deeper networks. After the forward-propagation of the samples, which just consists in passing the data through the network and calculating the loss at the output, the gradient is calculated back through the network, exploiting the chain rule of derivation. This technique is called backpropagation.

If the inputs have lots of dimensions, like images and time-series data, the weight matrices $W_{n\times k}$ can reach an unwieldy number of parameters. This kind of data usually presents strong local correlations, since pixels that are close together usually belong to the same natural feature, while the ones that are far apart might be uncorrelated. In standard neural networks implementations, every component of the input vector is connected to every component of the hidden layer. Connections between couples of pixels that are far apart constitute most of the total number of parameters while containing the least amount of information, so a common solution to this problem is to drop far apart connections by focusing on local relations and replacing the multiplication step in the neural network architecture with a matrix convolution. Each hidden layer presents different convolution kernels, usually a few pixels in size, and each one of them produces a "feature map" of the original image. Standard convolutions produce feature maps with the same size as the original image, but by implementing strategies such as max pooling or choosing higher pixel strides for the convolution, the size of the image will be shrunk, so that at the end the weight matrices will be reduced in size. These frameworks are called Convolutional Neural Networks (CNNs). An example of how a convolutional kernel is applied is present in fig. 3.1 Since this thesis will explore the movement of Superattenuators through 2-dimensional spectrograms, CNNs architectures will often be implemented.

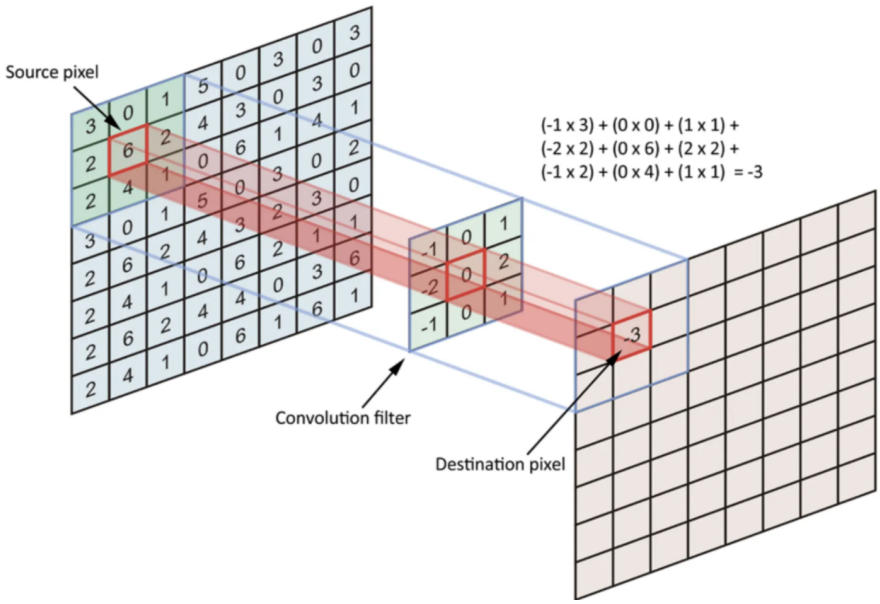

Figure 3.1: Example of how a Convolutional kernel is applied to 2D images. [40]

3.2 Machine Learning in Gravitational waves astronomy

Machine Learning tools have recently been finding their way into gravitational waves astronomy. At all of the stages of the pipeline, from detector noise characterization to parameter estimation of real gravitational wave signals, machine learning driven solutions are currently being explored. In many instances, these novel algorithms were found to equal or even surpass the analytical tools already put in place. A comprehensive review of applications of machine learning in gravitational waves astronomy can be found in the paper [41] and since this thesis is focused on how external noise influences the instrument, a summary of the current state of machine learning applications into the detector characterization pipelines will be the focus of this section.

As discussed in sec. 2.4, the noise affecting the interferometer is quite complex. External and instrument noises enter the detector strain $h(t)$ through linear and non-linear couplings, making the noise affecting the measurement both non-stationary and non-Gaussian. The vast amounts of data produced by the interferometer and the monitoring auxiliary channels can be an obstacle for manual analysis, but it can become a useful resource as a training dataset for machine learning algorithms. The short bursts of non-stationary noise called glitches and mentioned in sec. 2.4 are objects of multiple studies that take into consideration both the strain and the Auxiliary channel data. Projects like Gravity Spy [32] and GWitchHunters [33] use glitches data labeled by citizen scientists for the training of a supervised algorithm for the classification

of glitches. An algorithm with these capabilities operating in real-time can be useful both for rapidly excluding transient noises from gravitational waves searches and faster noise characterization.

Algorithms implementing information from auxiliary channels to characterize glitches have also seen some successes. The usual workflow consists of the use of datasets containing auxiliary channel data around glitch events. The presence of a glitch in $h(t)$ is used as a label for the auxiliary channels. Once trained, these algorithms should be able to predict the presence of a glitch in $h(t)$ just by looking at auxiliary channel data. A mature work in this direction is the iDQ algorithm [42], a supervised learning tool capable of rapidly vetoing the quality of the stream in $h(t)$. If a glitch is detected in the auxiliary channels non-sensible to gravitational waves, then the event is instantly excluded from gravitational wave searches. This has been proved useful in the rapid release of the data of GW170817, an event that happened in coincidence with a glitch in the LIGO-Livingston interferometer[43]. Another algorithm called EMU uses the full list of LIGO auxiliary channels to give a significant score to each of the channels concerning the production of the glitch.

Some works that characterize the ground motion at the interferometer's sites with machine learning have already been done. In the paper [44] a regression algorithm is trained on archival seismic data to infer the excess local ground motion at the different interferometer sites caused by distant loud earthquakes. The algorithm is also capable of predicting the impact on the measurement quality itself, and warnings issued by this algorithm have the possibility of alerting the operators of the interferometer to switch to a mode capable of withstanding the elevated ground motion without risking a lock loss [45].

Machine learning has also been used to directly clean the data output of the interferometer, by learning the non-linear coupling between auxiliary channels and $h(t)$, that standard analytical algorithms like Wiener filtering are not able to characterize. One of such algorithms [46] was able to successfully remove the non-stationary noise due to couplings with the power lines, for frequencies near the main at around 60 Hz.

3.3 Variational Autoencoders

The analysis of high-dimensional data, like images or time series data, can be quite challenging due to an effect called the "Curse of dimensionality". The volume where the data lives gets multiplied for each new dimension and exploration of this space becomes more and more computationally demanding, reaching impossible levels pretty fast. Datapoints sparsely populate this volume, and any distance measured between these points will be dominated by the noise in

the data. One possible way to overcome this issue is to characterize each data point with some summary statistics that should meaningfully represent the samples with much fewer parameters. Summary statistics to define general features typical of unmodelled signals could be the maximum loudness, the peak frequency, the duration, the bandwidth, and so on. By characterizing the signals with just these variables already a lot of different analyses could be easily run, but sometimes more information is needed to perform a more refined analysis. The knowledge of the underlying processes that originated these signals could give hints to find other useful summary statistics, but not always these processes are known. This is where machine learning driven generative models get into the picture. By training a machine learning framework to recreate these signals starting from some latent variables z, the algorithm will need to build an internal model for the generation of the data based on z. These variables can now be used as summary statistics that define each sample. One of these algorithms is the Variational Autoencoder.

Autoencoders are machine learning tools capable of learning a generative model for arbitrary data distribution. These frameworks are forced to find the relevant information contained in the data by making it pass through an "information bottleneck". Given a dataset containing samples $\mathbf{x} \in R^n$, the framework passes the data through a parametrized function $f(\mathbf{x})_\theta : R^n \times R^m$ called the encoder, which compresses the inputs x_i coming from the original high dimensional space down to a lower dimensional one $\mathrm{z}_i \in R^m$ where $m < n$. The space R^m is called the latent space. Another function $g(\mathrm{z})_\phi : R^m \times R^n$ decompresses the data back to its original size by creating a sample $\tilde{x}$. The framework is then trained with an optimization problem to minimize the distance between x and $\tilde{x}$, through some user-defined norm, the most common of which is the Mean Squared Error (MSE). If $f(\mathrm{x})_\theta$ and $g(\mathrm{y})_\phi$ are linear functions the framework will collapse to a Principal Component Analysis, but if non-linear activation functions are implemented, more complex behaviors can arise. The training loss function for a standard Autoencoder trained on a dataset with N samples and MSE as a distance norm will be:

$$L(\theta, \phi; x_i) = \frac{1}{N} \sum_{i=1}^{N} \|x_i - g(f(x_i; \theta); \phi)\|^2 \tag{3.5}$$

In practice, naif implementations of Autoencoders tend to create quite irregular latent spaces, resulting in points close in their latent representation being quite different in the visible units. This defies our intuition of summary statistics, and so the need to put some kind of regularization of the latent space arises. A possible solution is presented by a modification of the standard autoencoder proposed in [47] Called the Variational Autoencoder (VAE). In VAEs the mapping of the encoder is probabilistic, where each sample is encoded into a probability distribution over the latent space $p_\theta(\mathbf{z}|\mathbf{x})$, which, in Bayesian terms, is considered the posterior distribution. The regularization is then obtained by imposing a prior $p(\mathbf{z})$ over the latent space, which "pulls together" all the different sample distributions towards a common one. The qualita-

tive effects are that now the model is incentivized to partially "overlap" the $p(\mathbf{z}|\mathbf{x})$ distributions, finding common features between different samples and coding their behavior on a specific latent variable. For example, VAEs trained on datasets of faces tend to encode in some specific latent variables features like the hair color, the "happiness" of the expression and the inclination with respect to the camera [48].

In almost every use, the probability distribution of choice for VAEs is a multivariate Gaussian $\mathcal{N}(\mu, \Sigma)$, parametrize d by the mean vector μ and the covariance matrix Σ, which for convenience purposes is usually kept diagonal. The prior will be the standard Gaussian distribution $\mathcal{N}(0, \mathbb{I})$ where $\mathbb{I}$ is the identity matrix. In actual implementations, the encoder function $f(x; \theta)$ will deterministically produce both a means vector μ and a variance vector σ. For each training epoch, the posterior of each training data point will be sampled from $\mathbf{z}_i \sim p_\theta(\mathbf{z}|\mathbf{x}_i)$ to produce the reconstructed outputs. The loss function is then calculated by measuring the distance between input and output. Normally it is not possible to backpropagate the gradient through the random sampling process, but the reparametrization of the Gaussian sampling makes this possible. The actual output $\mathbf{z}$ will be equal to

$$\mathbf{z} = \sigma_\theta(\mathbf{x})\zeta + \mu_\theta(\mathbf{x}) \quad \zeta \sim \mathcal{N}(0, \mathbb{I}) \tag{3.6}$$

The regularization of the latent space is obtained by calculating the Kullbach-Liebler (KL) divergence between the posterior and the prior. Another reason that makes the Gaussian the probability distribution of choice of VAEs is the fact that the KL divergence can be easly computed analytically. For this specific implementations, where the distance is between a multivariate Gaussian with a diagonal covariance matrix and a standard Gaussian, the KL divergence will be calculated by:

$$\mathcal{D}_{KL}(p_\theta(\mathbf{z}|\mathbf{x_i})\|\mathcal{N}(0, \mathbb{I})) = -\frac{1}{2}\sum_{k=1}^{K}(1 + \log\sigma_k^2 - \mu_k^2 - \sigma_k^2) \tag{3.7}$$

And finally, the total loss function used to train the VAE will be:

$$\mathcal{L}(\mathbf{x}; \theta, \phi) = \frac{1}{N}\sum_{i=1}^{N}\|x_i - g_\phi(\mathbf{z} \sim p_\theta(\mathbf{z}|\mathbf{x_i}))\|^2 + \mathcal{D}_{KL}(p_\theta(\mathbf{z}|\mathbf{x_i})\|\mathcal{N}(0, \mathbb{I})) \tag{3.8}$$

And in fig. 3.2 a schematics of the VAE architecture is present.

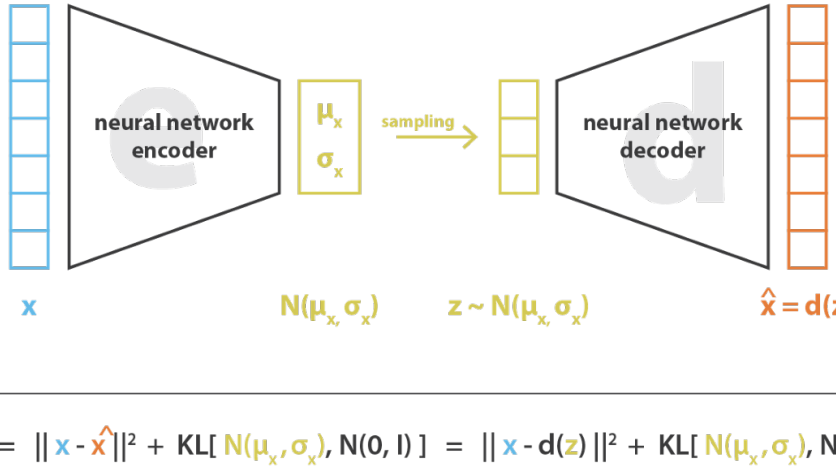

Figure 3.2: General overview of a VAE architecture, similar to the one that will be used in the final project. [49]

The samples encoded in the latent space will have dropped all of the redundant information and most of the random noise that comes with the original Data. This makes VAEs a great tool for denoising and since the latent space will have only the most relevant information this will present the clustering algorithm with only the main morphology of the images and not random noise.

3.4 Clustering with Gaussian Mixture Model

Now that the data has been cleaned of noise and contains only relevant information, a pattern-finding algorithm can be deployed. The use of VAE in conjunction with clustering algorithms has proven successful in previous works with highly dimensional datasets [50] and so this thesis will try to reproduce those results on data coming from the Superattenuators of Advanced Virgo. A few clustering algorithms have been tried, but the best-performing one seems to be the Gaussian Mixtures Model (GMM).

Like VAE, the Gaussian Mixtures Model can be considered a generative algorithm. It takes the assumption that the data points were sampled from a probability distribution that takes the form of a mixture of K Gaussian $\mathcal{N}(\mathbf{x}|\mu, \mathbf{\Sigma})$. Each Gaussian has its parameters μ_k and Σ_k and its weight in the final probability distribution π_k. The samples $\mathbf{x}$ will be considered to be drawn from

$$p(\mathbf{x}|\mu_k, \Sigma_k, \pi_k) = \sum_{k=1}^{K} \mathcal{N}(\mathbf{x}|\mu_\mathbf{k}, \mathbf{\Sigma_k})\pi_k \tag{3.9}$$

The likelihood of observing a certain realization of the dataset $\mathbf{X} = \{x_1, x_2 ... x_N\}$ will be:

$$p(\mathbf{X}|\mu_k, \Sigma_k, \pi_k) = \prod_{i=1}^{N} p(\mathbf{x_i}|\mu_k, \Sigma_k, \pi_k) \tag{3.10}$$

To find the parameters of the GMM that better represent the data, an optimization algorithm is run to maximize this likelihood and to find the probability distribution that best fits the data.

Now that the total probability distribution has been found, one can turn this into a clustering algorithm by assigning each sample to the Gaussian k that most likely generated it. By assigning a vector $\mathbf{z} \in R^k$ to each sample, which will take the value of 1 if the sample $\mathbf{x}$ belongs to the k-th cluster and 0 otherwise, the latent variable needed to consider this a generative model is set. Since now the GMMs is a generative model, the probability of having a sample x as a realization of the latent variable z can be defined as

$$p(\mathbf{x}|\mathbf{z}; \mu_k, \Sigma_k) = \prod_{k=1}^{K} \mathcal{N}(\mathbf{x}|\mu_k, \Sigma_k)^{z_k} \tag{3.11}$$

and by using Bayes theorem one could invert this quantity and get the probability that a certain sample x_i comes from the k-th cluster with

$$p(z_k = 1|x; \mu_k, \Sigma_k, \pi_k) = \frac{\pi_k \mathcal{N}(\mathbf{x}|\mu_k, \Sigma_k)}{\sum_{j=1}^{K} \pi_j \mathcal{N}(\mathbf{x}|\mu_j, \Sigma_j)} \tag{3.12}$$

Summary

This chapter had the objective of giving an overview of the machine learning tools used in this project. First, the principles behind machine learning were given, with a focus on unsupervised learning, since it is the main framework used in the thesis. Then an overview of the previous works that used machine learning in gravitational waves astronomy, in particular how it was used to characterize the detector was given. For the rest of the chapter, the specific algorithms used in the thesis were described, so the VAE for the dimensionality reduction and the GMM for the clustering.

Chapter 4

Studying the Superattenuators with machine learning

In this chapter, everything introduced in the previous parts of the thesis will be put together to construct an algorithm capable of clustering together the responses of the Superattenuators in periods of high seismic activity. Standard analytical tools make it hard to study such a complex phenomenon, so the algorithm used in this thesis blends standard analytical tools, deep generative models, clustering algorithms and Bayesian statistics to reach its objective. This chapter will open with sec. 4.1 which will be a description of the pipeline that was used to construct the dataset. In sec. 4.2 the seismic landscape at Virgo will be analyzed to design an algorithm capable of detecting periods of high seismic activity automatically, which will be described in sec. 4.3. The way the Superattenuator's time-series are turned into spectrograms is described in sec. 4.4 and sec. 4.5 will analyze the dataset as a whole. Finally, Sec. 4.6 will describe the architecture of the dimensionality reduction algorithm.

4.1 Data acquisition pipeline overview

To download and pre-process the data for the clustering algorithm, a specific pipeline had to be put in place. Many things needed to be taken into account while designing this tool, since machine learning analysis focused on Superattenuator's data has never been tried before.

The work started by choosing a period of interest to run the analysis on. The project focuses on the activity for 3 weeks from 29 April 2019 to 20 May 2019. This period has been chosen since it presents quite a varied mix of conditions under which the interferometer is operating. Some days present a high number of low-frequency glitches, as signaled in the Gravity Spy Virgo dataset [32]. Other days present no glitch activity at all, indicating good data quality periods and others again have been flagged for bad data. This time-span also cointains the Fig. 4.1

reports a plot of the evolution of the glitch count as well as the interferometer index number, that encapsulates what state the detector is in.

After the period had been chosen the analysis of the seismic noise registered by the seismometers present at the 3 main Virgo buildings was conducted.

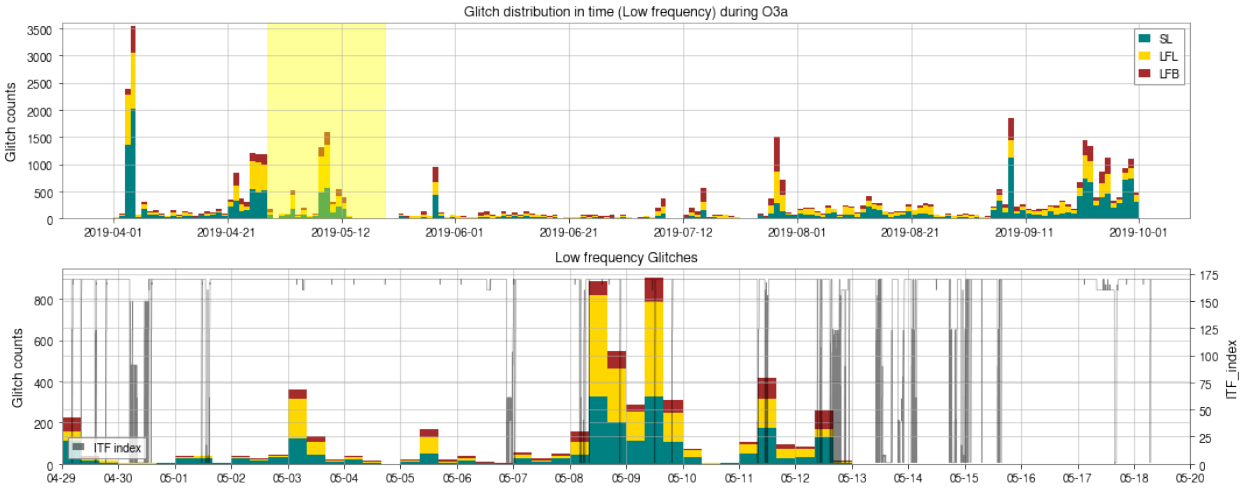

Figure 4.1: Distribution of low frequency glitches during O3a(top) and specifically during the period of the analysis (bottom) as well as the interferometer lock index. The glitches plotted are Scattered Light (SL) Low-Frequency Lines (LFL) and Low-Frequency Bursts (LFB).

The results of this analysis were used to design the actual algorithm for data download. This algorithm, based on the calculation of the Band Limited Root Mean Squared error of the seismic channels, informs at what time and in which building the seismic activity is higher than normal, indicating that probably the Superattenuators are reacting to it.
The channels downloaded are the 4 DoFs that define the F0 position of each Superattenuator present in the building where the seismic excitation is. To better highlight frequency features, spectrograms are obtained from the time-series to represent the data.
Once all of the 4-channel 2D spectrograms are prepared, the dimensionality reduction algorithm VAE can construct a generative model for the dataset. That will help to find summary statistics that characterize each sample.
These compressed samples are now of a low-enough dimension to run a clustering algorithm on them. Now these clusters can be explored and their properties examined, to see if some insight on how the Superattenuator behaves and influence the quality of measurement can be inferred

4.2 Seismic landscape at Virgo

Since the main goal set for this thesis regards how the mirror suspensions react to seismic excitations, the first step needed to develop an algorithm capable of detecting these events is to first understand the seismic landscape at the Virgo detector. The ground motion at the site is monitored by Triaxial Guralp 40T seismographs [51] that are placed in every main building at the site (CEB, WEB and NEB), providing 3-dimensional information on the ground motion, in the North, West and Vertical directions. These instruments provide 500Hz data all year round and are active even when the interferometer is not in science mode. The typical power spectral density for the ground motion of all of the channels is present in fig. 4.2. This distribution tends to change over time, and so to better characterize this evolution in the different frequency ranges, the Band Limited Root Mean Squared error (BLRMS) was used. If $S^{1/2}(f)_{[t,t+\Delta t]}$ is the PSD of a signal, obtained for the interval $[t, t+\Delta t]$ The BLRMS will be obtained by integrating this function in a certain frequency range $[f_1, f_2]$

$$\mathrm{BLRMS}(f_1, f_2)_{[t,t+\Delta t]} = \int_{f_2}^{f_1} S^{1/2}(f)_{[t,t+\Delta t]} df \tag{4.1}$$

This quantity represents the power that the signal contains in that specific frequency range.

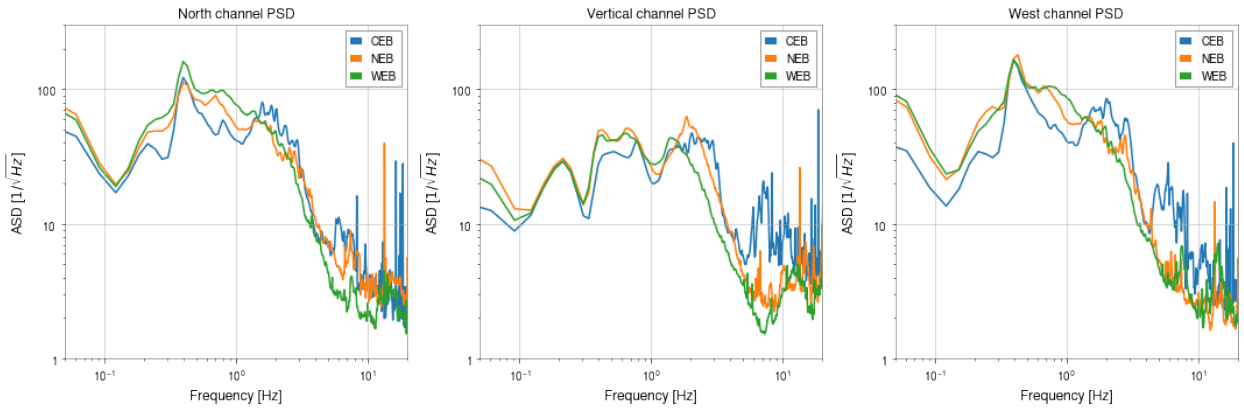

Figure 4.2: PSD of seismic activity calculated during 2 May 2019 for the different EGO buildings and directions.

The peculiar shape and evolution of the microseismic PSD arises from the sum of different noise sources acting on different parts of the spectrum. Going up the frequency range, these sources will now be explored, following roughly the findings in [45].

- **Sea activity**: these are the disturbances generated from waves in the Tyrrhenian sea interacting with the coast. They produce seismic waves with frequencies comparable with

the sea waves, between 0.1 Hz and 1Hz, peaking at 0.35 Hz and they represent the most prevalent source of microseismic activity at Virgo. It is known that this is the case in many parts of the world [52] and especially at the Cascina site since the coast is only 15 km away. Also, Periods of high sea activity have been linked to the general ground motion at the site. This noise source is strongly seasonal, with high peaks in the winter months

- **Wind activity**: High winds have been shown to impact the quality of the detection [45] and also the ground motion at the site. Wind intensity couples to ground motion in many different ways: high wind periods can contribute to high sea activity, hightening the sea microseismic peak; wind can cause the buildings to shake, and induce elevated activity of the Pontedera wind farm, positioned 7 km away from the site. This has been shown to be responsible for a microseismic peak at Virgo coinciding with the turbine typical rotation frequency of 1.7 Hz[53]. The wind speed and direction will be part of the analysis. The wind mostly blows in the East-West direction as can be seen in fig. 4.3, usually inverting direction between the day and the night.

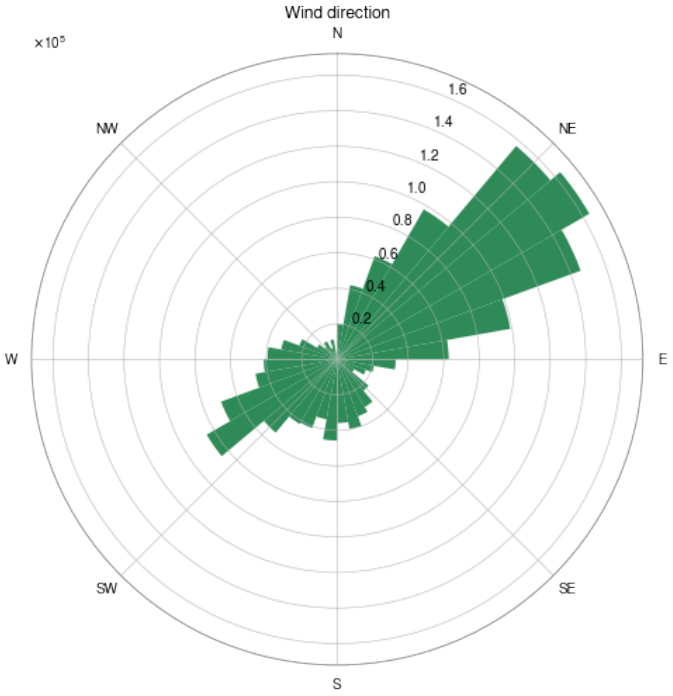

Figure 4.3: Histogram of the wind direction at the Virgo site, calculated over the 3 weeks that the analysis took place.

- **Antropogenic activity**: above 1Hz anthropogenic activity dominates the spectrum. The BLRMS in the region 1-4Hz has been shown to follow a day and night cycle, as well as a weekly one, being quiet at night and during the holidays while being loud during daytime

and working days. This behaviour is quite evident when looking at fig. 4.4 and fig. 4.5. The main culprits are thought to be road activities and more specifically the bridges of the FI-PI-LI freeway [54], about 3 km away from the main building. The peaks in the seismic activity pretty much correlate with the resonant mode of the bridge pillars [55].

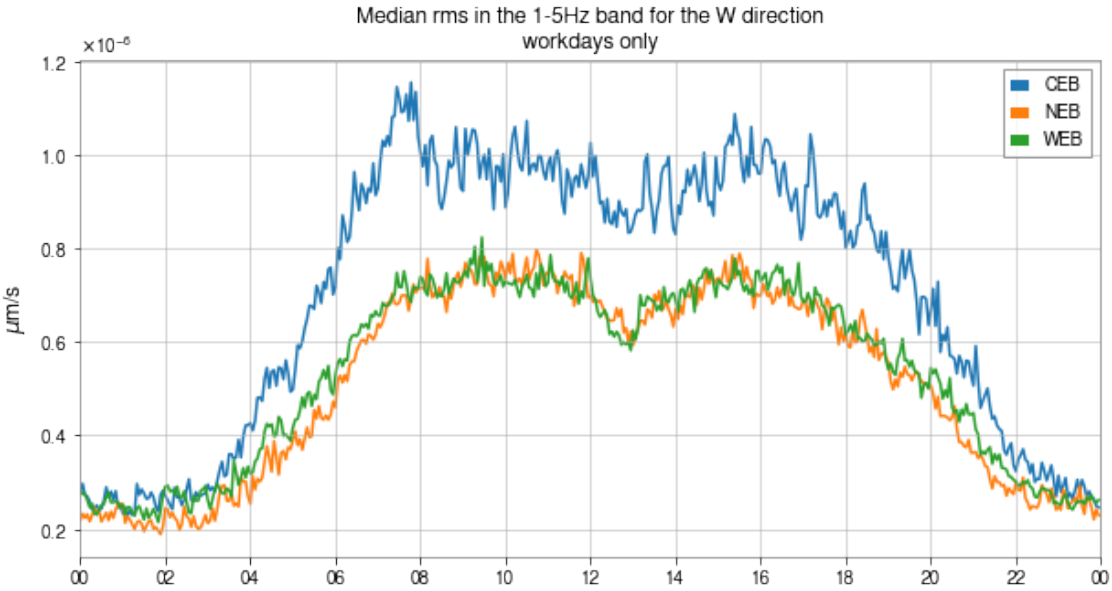

Figure 4.4: Evolution of the BLRMS for the 3 buildings in the 1-5 Hz band averaged over the workdays of the 3 week of the analysis. The daily modulation correlating with the antrophic activity is quite clear.

- **Local sources**: As the frequencies increase, ground oscillations get more and more dampened[45]. The only noises above 10 Hz that survive come from local disturbances, happening close to or even inside the buildings and which are usually transient in nature. These sources of noise can include the movement of vehicles operated by the EGO staff, farming activity happening around the site as well as fans and compressors inside the buildings. A predictable source of noise comes from the crew operations happening each Tuesday morning, where heavy trucks move near the buildings to refill the liquid Nitrogen tanks of the cryopumps needed to maintain the vacuum inside the detector arms [45].

In fig. 4.5 the weekly evolution of the BLRMS in 3 different bands, corresponding roughly to the Sea activity, anthropic noise ad local activities frequency bands are present. The anthropic origin of the two latter bands is pretty evident, given how closely they follow the human work cycle.

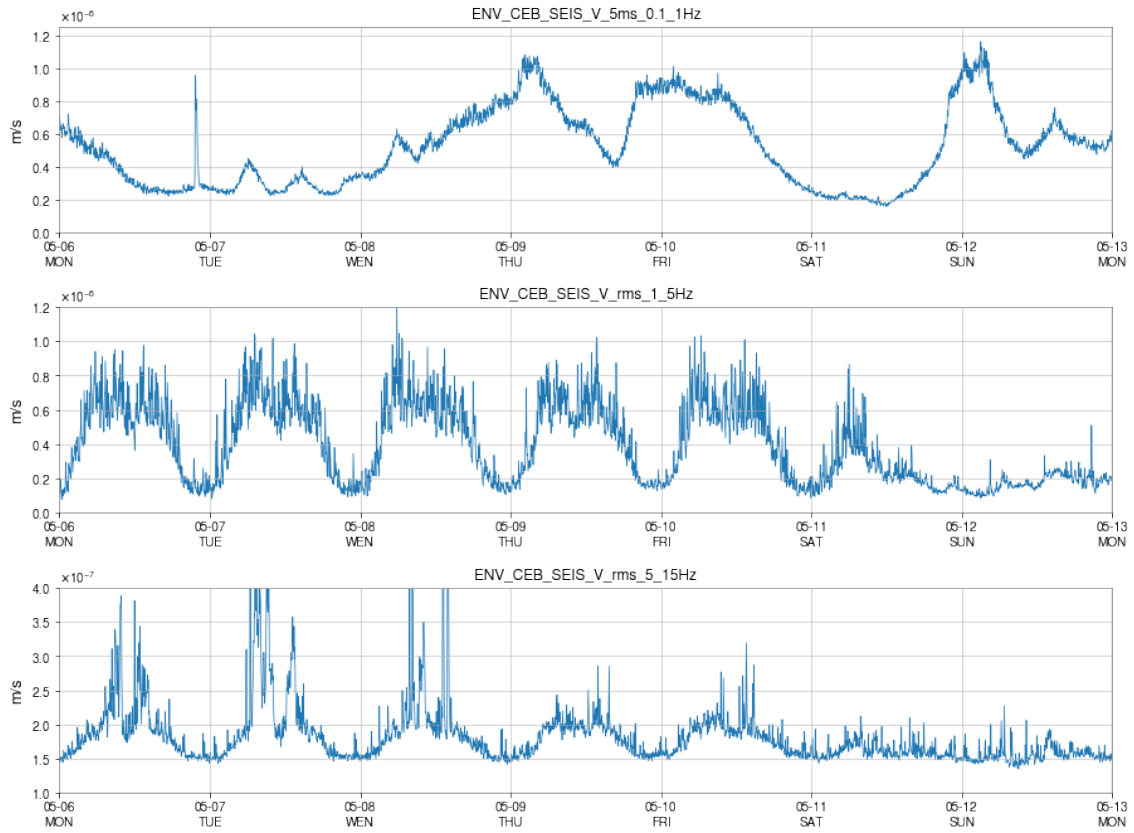

Figure 4.5: Evolution of the BLRMS over a week during O3a. The bands $[0.1, 1]$Hz, where the effects of the sea activity are the most prominent, and they do not show any correlation with the weekly activities, the $[1, 5]$Hz band, correlated with far away vehicles movements and the $[5, 15]$Hz band mostly linked to local activities. For both of these bands, the link with the workweek is pretty clear, and also it is evident how much more variability is present in the local band.

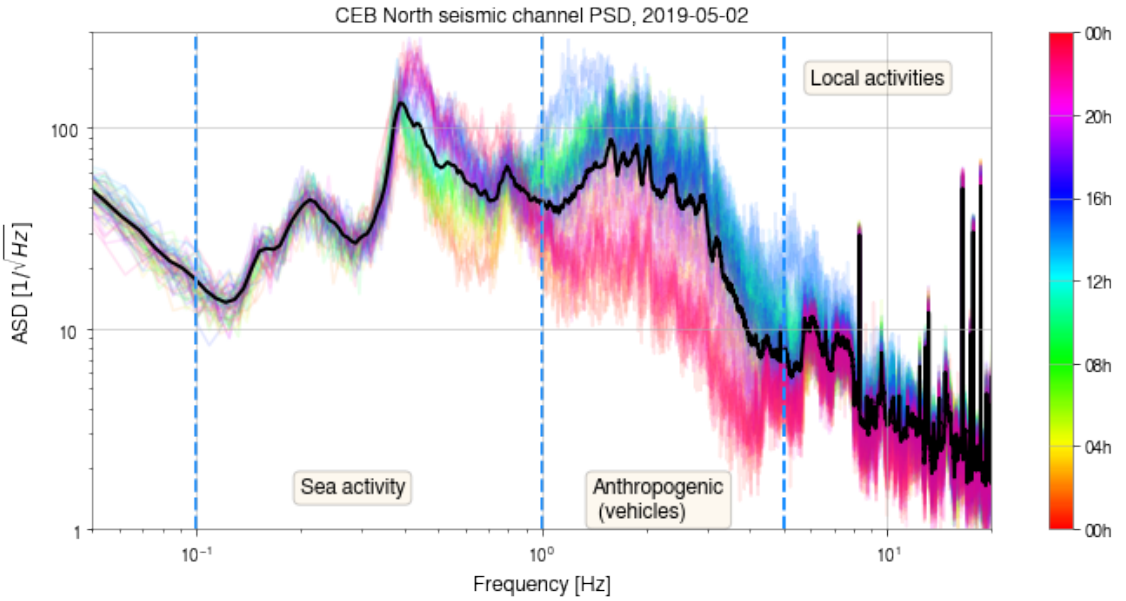

Figure 4.6: mean PSD of the North channel of CEB during the 2 of May 2019. The black line presents the median, while the colored lines are the single 1280 seconds segments from which it was calculated. The color of the segments depends on the time of day it was taken, to represent how the seismic activity varies over the day.

4.3 Data Acquisition

The Seismic landscape at the Virgo interferometer is quite heterogeneous, varying widely from building to building, day to day, season to season and even hour to hour. By observing the PSD is also evident how the intensity spans different orders of magnitude over the whole frequency range. To develop an algorithm capable of detecting seismic excitation that will have a noticeable effect on the Superattenuators, all of these factors had to be taken into consideration.

I designed a pipeline to download the Superattenuator data corresponding to elevated periods of seismic activity. The data was obtained by accessing the records of the auxiliary channels of the interferometer. The historical data is stored locally at Virgo in h5d5 format, accessible via .ffl files. The files have been downloaded with the help of the python library GWdama [56], which aides in the handling of these kinds of files. The channels for the monitoring of seismic triggers chosen are the 3 cardinal directions ((N)orth, (W)est and (V)ertical) for the 3 main buildings of Advanced Virgo (CEB,NEB and WEB). The data stream for these channels has a sampling frequency of 1000 Hz but for this project, it was downsampled to 100 Hz.

The algorithm that detects periods of high seismic activity is based on the calculation of the

BLRMS, to have an idea of the power cointained in the frequency bands that the Superattenuators are more sensible to. Due to the variability of the seismic intensity in time and in place, it was chosen to avoid a method based on a fixed threshold trigger. Instead, the method calculates a running average of the BLRMS and takes this value as the floor power. For this project, it was decided to take one hour of data around each timestep and divide it into $\Delta t = 16$ seconds segments to calculate the average. To actually trigger the algorithm the ratio between the BLRMS of each segment and the running average is calculated. Whenever this value surpasses a certain user-defined threshold, the download of the Superattenuator data is triggered. After a bit of experimentation, the best threshold value was found to be 4. Lower threshold values would lead to too many triggers, with many of them exhibiting unclear or totally absent features in the Superattenuaotr spectrograms, while higher values would lead to too few data points, not enough to train the algorithm on.

This algorithm taken as a whole can be seen as a way to look for 4σ significant events happening in the specific band of interest. For the i-th interval $[t + i\Delta t, t + i\Delta t + \Delta t]$ the equation used to determine the trigger will be

$$\frac{\mathrm{BLRMS}(f_1, f_2)_i}{\frac{1}{225}\sum_{j=-112}^{112} \mathrm{BLRMS}(f_1, f_2)_{i+j}} > 4 \tag{4.2}$$

The limits of the sum are defined this way since there are 225 windows of 16 seconds in an hour. This process is run for each building and each one of the 3 DoFs. If the threshold is surpassed in any of the 9 channels, the download of the data for each Superattenuator inside the trigger building is initiated.

For the Superattenuator data, a 64-second window was chosen to be downloaded around the trigger. The channels downloaded are the x, y, z and θ_y position of the F0, which define the suspension point of the Superattenuator pendulum chain. As mentioned before, y is the vertical component, while x and z define the horizontal plane, with z being oriented as the direction of the laser. This data is obtained by LVDT (Linear Variable Differential Transformer) sensors mounted on the top ring, which measure the distance between the filter and the walls of the vacuum enclosure. These instruments have a sampling frequency of 10kHz, but for storage purposes, these have been reduced in the files stored in the Virgo data center to 500 Hz. I performed another downsampling to 100 Hz for better data storage. As it is the case with the seismometer, these channels are always available, even when the interferometer is not in science mode.

Since the Superattenuators are tuned to dampen high frequency components of ground motion, the project will be focused on the low part of the spectrum. Three frequency bands were taken as study goal for this project. $[0, 0.5]$Hz , $[0.5, 5Hz]$ and $[5, 10]$ Hz. If the BLRMS

algorithm shows an excess of power in any of the 3 bands, the download is triggered.

4.4 From time series to frequency space

In previous works that implement machine learning algorithms in gravitational wave astronomy, like glitches classification, [32] and auxiliary channel characterization [42], spectrograms are often chosen over raw time series as a way to present data to the neural networks. The main reason is due to the large sensitive frequency range of the detector (10-2000 Hz for Virgo) and the fact that the gravitational signal spans good portions of it. This implies that really different scales need to be taken into consideration in the time series representation, while, when looking at spectrograms, all of these features are presented at similar scales at the same time, making it easier for both humans and machines to study these features.

Usual algorithms for computing spectrograms, like straight FFTs, were explored for this work, but these have quite a few limitations when exploring a vast frequency range. The frequency resolution Δf and the time resolution Δt are related to each other like $\Delta f \propto 1/\Delta t$: to have good resolution at low frequencies one would need quite a long time window, but this comes at the price of sacrificing time resolution on the higher frequency region, where events usually happen at shorter timescales. One possible solution that can be implemented is to have a window that adapts its length based on the frequency it is calculating. This is the strategy adopted by the Q-transform algorithm, which is quite extensively used both in music spectrograms visualizations as well as in gravitational wave astronomy [57].

4.4.1 Q-transform

The Q-transform is a type of wavelet transform, where, contrary to the usual Fourier transform, the base of infinite sines is replaced with Gaussian sine wavelets, limited in both time and frequency domains. This is obtained by adding a window therm on the time domain, and in the Q-transform this gets smaller as the frequencies increase. The Q-transform of a time series $x(t)$ will be equal to

$$X(\tau, \phi, Q) = \int_{-\infty}^{+\infty} x(t) w(t - \tau, \phi, Q) e^{-2i\pi\phi t} dt \tag{4.3}$$

$X(\tau, \phi, q)$ is the average amplitude and phase of a signal around a region around time τ and frequency ϕ, while w is the window function that depends on the quality factor Q. The window chosen is a Gaussian and will have the equation:

$$w(t - \tau, \phi, Q) = \frac{W_g}{\sigma_t \sqrt{2\pi}} \exp\left[-\frac{1}{2\sigma_t^2}(t - \tau)^2\right] \tag{4.4}$$

Where W_g is a normalization factor and $\sigma_t = \frac{Q^2}{8\pi\phi^2}$. The window defines also the resolution in frequency space, and since the Fourier transform of the window function will still be a real-valued Gaussian, it is possible to know its variance $\sigma_f^2 = \frac{2\phi}{Q^2}$. While in frequency space σ_f is directly proportional to ϕ, in the time domain $\sigma_t \propto 1/\phi$. This indicates that for higher frequencies, there will be a greater time resolution with respect to the lower frequencies, making it possible to characterize better the short transient features that are more likely to happen in this region. In turn, lower frequencies will have higher resolutions in the frequency domain. Since these two properties hold at the same time for a single spectrogram, this is the reason why Q-transforms are so useful in studying signals that preset broadband features.

Another important step left to produce a representative spectrogram is to have a good choice of the quality factor Q. This determines the "trade-off" between frequency and time resolution. Lower values of Q will lead to better characterization of features with fast time evolution at the cost of a bad frequency resolution, and vice-versa. In the actual implementations, this is chosen automatically by the algorithm by finding the quality factor that presents the maximum value in the spectrogram output. How different values of Q affect the representation of the glitches in $h(t)$ is present in fig. 4.7.

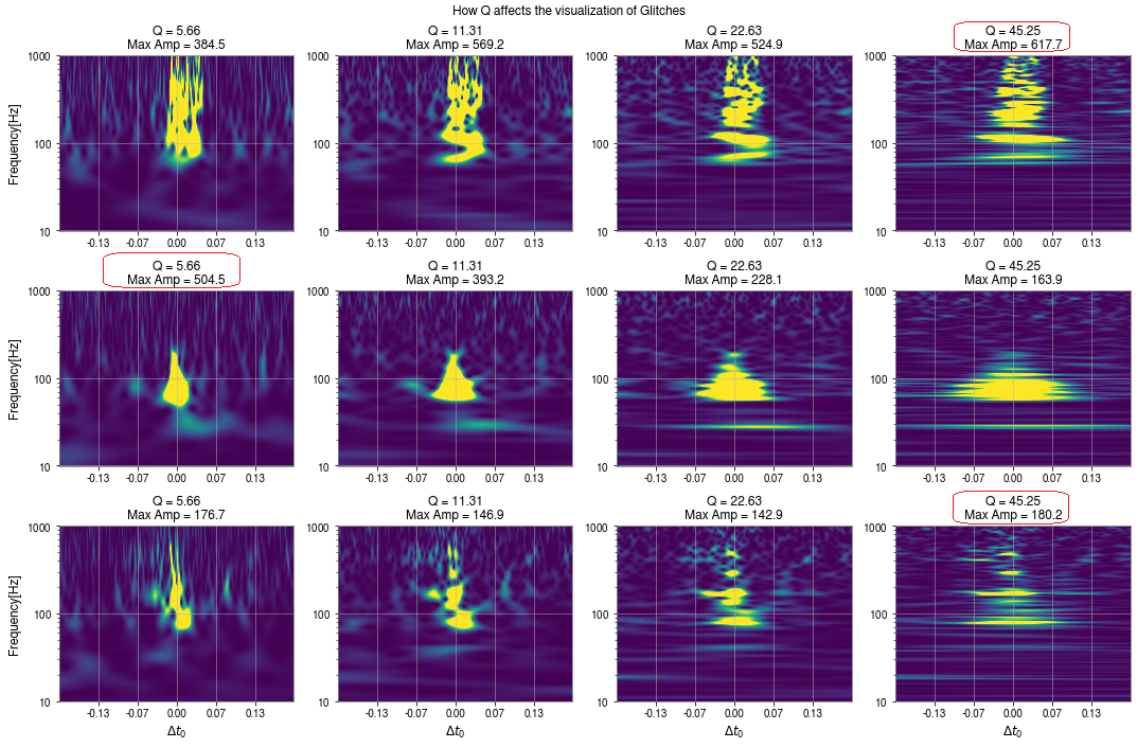

Figure 4.7: Three glitches of the Virgo interferometer represented with different values of Q. It is evident how the algorithm trades the frequency and the time resolution as this value increases. Highlighted in red are the highest values reached in the spectrograms, which represent which Q the algorithm deems is the best.

The implementation of the algorithm uses a finite approximation of the Gaussian window for computational purposes. These limits are then used to define tiles of the time-frequency-Q space, with width dependent on the σ_t and σ_f values. An example of the tiling of the algorithm is present in fig. 4.8 [57]. The finite nature of the windows also implies that a few anti-aliasing precautions can be taken, in particular

$$Q \geq \sqrt{11}, \quad \phi \leq \frac{f_{\text{nyquist}}}{1+\sqrt{11}/Q} \tag{4.5}$$

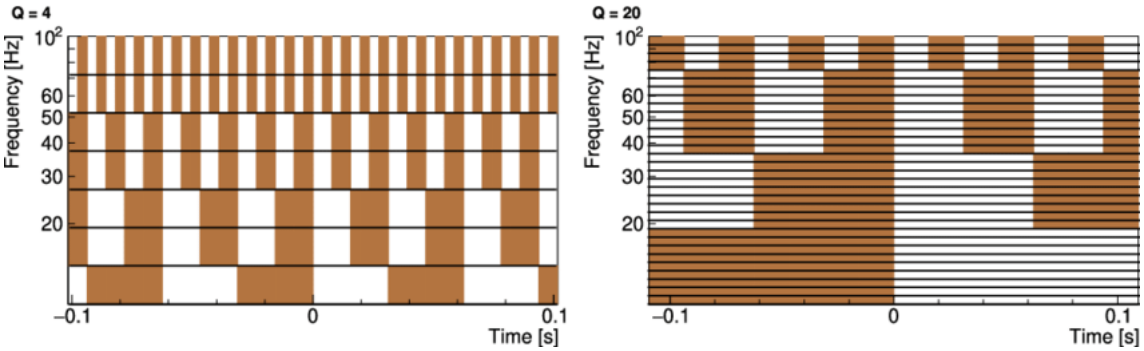

Figure 4.8: Tiling structure of the Q transform for different values of Q [57]

This indicates that the maximum frequency allowable by the Q-transform is a bit lower than the Nyquist frequency.

The Q-transforms implemented in this project, for the 4 F0 LVDT channels, are 20 seconds long and centered around the Superattenuator's spectrograms peaks. The frequency range of interest has been chosen to be between 0.2 and 20 Hz, while the total resolution is set to be 128x256 pixels [frequency x time]. To allow the algorithm to make correct tiles of the whole spectrogram space the quality factor was set to be in the range between 4 and 16 since otherwise, the anti-aliasing requirements would cut some of the frequency of interest.

Machine learning algorithms tend to perform best when the output they are trying to predict is limited to a certain range [36], so the last step needed to pass these spectrograms to the network is to bound them in intensity. An upper saturation threshold of 15 and a lower saturation at 0 on the Q-transform output are put in place. The lower bound is needed since the interpolation algorithms used to smooth the Q-transforms might output negative values, even if these are un-physical. The value of 15 was found after a few trials and errors. Setting a threshold this low might obscure some of the details of the louder features, but it preserves the main morphology of the image, making it more clear. Finally, the spectrograms values are dived by 15 to make the whole dynamic range fit inside the $[0, 1]$ interval.

In fig. 4.9 there is a summary of how the whole dataset generation algorithm flows.

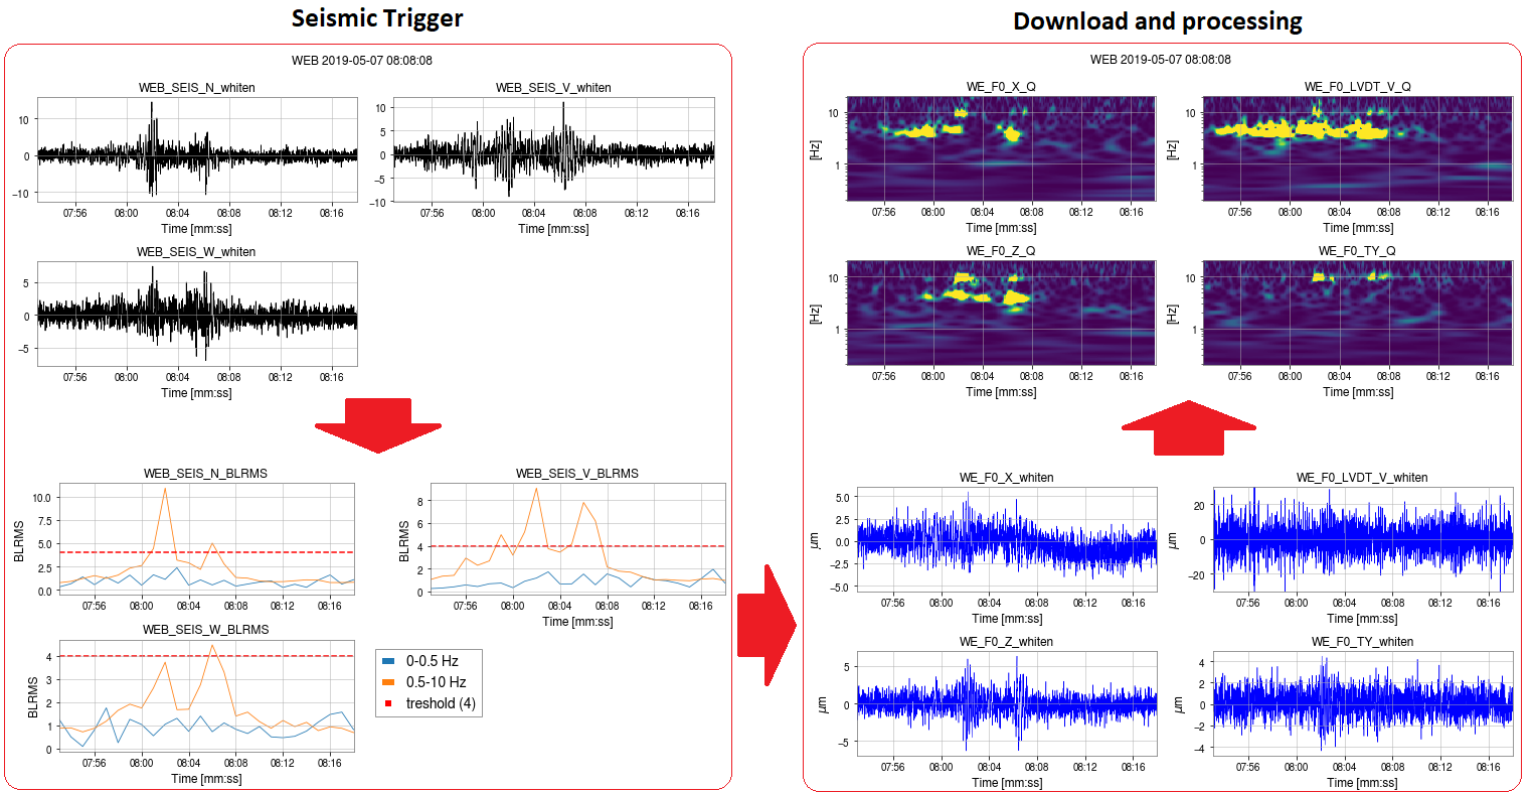

Figure 4.9: Schematic representation of the pipeline used in the project. The pipeline starts with the analysis of the seismic data on the upper left. the data has been whitened for visualization purposes. Then the BLRMS is calculated for each time-step, for this visualization reduced to 1-second windows. Then the corresponding 4 channels of the SA are acquired, and finally, the Q transform can be computed

4.5 Dataset statistics

In the 3 weeks of the study, a total of around 7000 seismic events were flagged by the algorithm, resulting in the download of around 13000 data samples of Superattenuators. This section will provide a broad analysis of this dataset. This is a fundamental step in the process, both for catching potential biases intrinsic to the dataset and to present an overview of the channels that will be used to analyze the clusters. The properties that will be looked at when analyzing the clusters are:

- The actual Superattenuator that originated the spectrogram
- Which channel presents the loudest signal, both for the F0 data and for the seismometers
- Peak Superattenuator and seismic frequencies of the spectrograms
- Binary Neutron Star (BNS) range, which is a measure of how sensitive the interferometer is at a given time
- Interferometer Lock Number, to monitor in which state the interferometer was at the time of the event
- Time and Weekday of the events;
- Wind speed and direction at the site ;
- Possible correlations with Glitches in $h(t)$

Fig. 4.10 reports the building-wise distribution for the seismic event, as well as the Superattenuators distribution in the final dataset. Most of the seismic activity seems to be involving the WEB and NEB buildings, while the CEB only accounts for $\sim$ 1000 events. Still, Superattenuators from the CEB represent around half of the dataset, since each seismic event that involves the CEB triggers the simultaneous acquisition of data from 8 different Superattenuators. The main culprit of this disparity is tough to be the higher noise floor originating from the local activities and the amount of equipment present at CEB. This creates a higher background noise when compared to the quieter end buildings, especially in the highest parts of the studied frequency range. This can be clearly seen in the vertical channel of the CEB in fig. 4.2. The higher noise floor in the CEB implies that the energy required to have a 4σ event is much higher than in the terminal buildings, making these events quite rare. This high noise floor is thought to also be the culprit behind the different daily distributions of seismic activity between the central and the end buildings, reported in fig. 4.11. The NEB and WEB seismic activity seems to be correlated with the anthropic day-night and weekly cycle, while for the CEB the cycle is reversed. As said

before, the high noise floor present during the day and pictured in fig. 4.6 make 4σ events rarer.

It is important to note that, since the analysis only considered at 3 weeks of data, statistics regarding the day of the week that events took place will be quite noisy. Also, since the 1st of May holiday is within the studied period and happened on a Wednesday, data regarding this weekday have to take this into account in order to consider effects of anthropic activities.

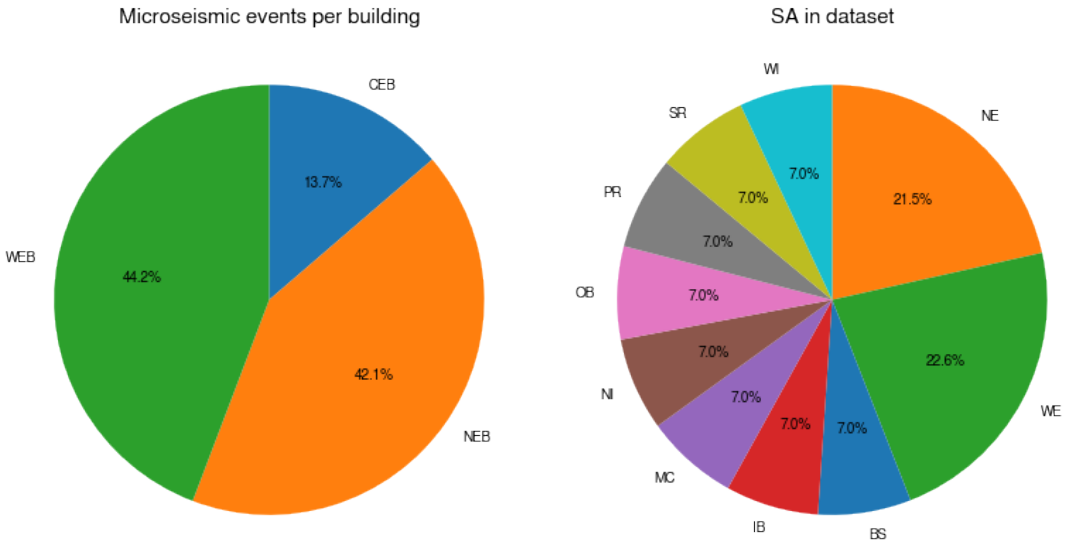

Figure 4.10: Distribution of the origin of the seismic activity identified by my algorithm. On the left the building of origin is presented while on the right the distribution of the single Superattenuators is present

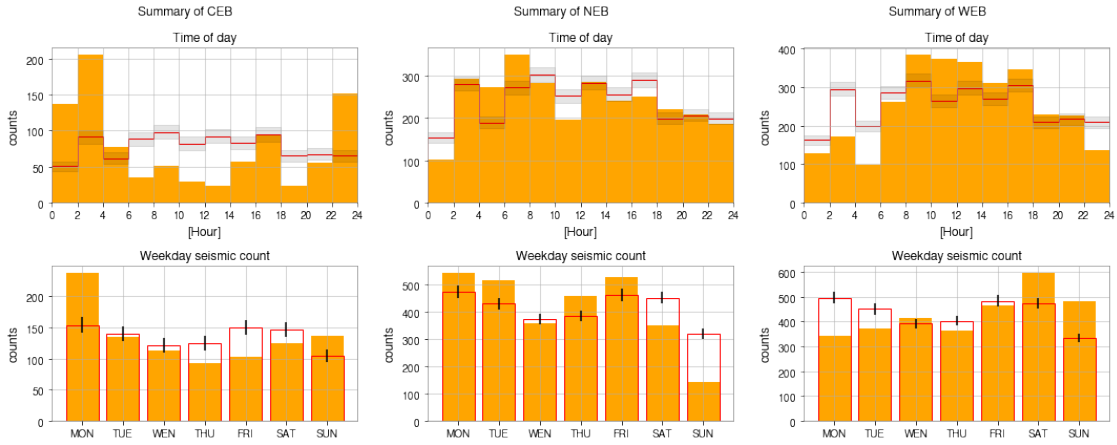

Figure 4.11: Daily and weekly distribution of events coming from the different buildings. It is clear how the day-night cycle affects the number of events in each of the buildings.

In fig. 4.12 the time distribution of small earthquakes, and their relation with the interferometer status is plotted. There seems to be no strong correlation between the seismic count and the glitch count, but there is a clear factor linking the most violent seismic "swarms" and periods of bad science and lock loss of the interferometer. Periods of elevated ground motion are known to cause unlocks of the interferometer and are usually caused by the seismic waves of distant loud earthquakes resonating around the planet. In particular, the peak of seismic activity happening on the night between the 6 and 7 of May seems to be linked to a 7.2 magnitude earthquake happening in Indonesia, and its effect on the interferometer has been the subject of a study presented in [45].

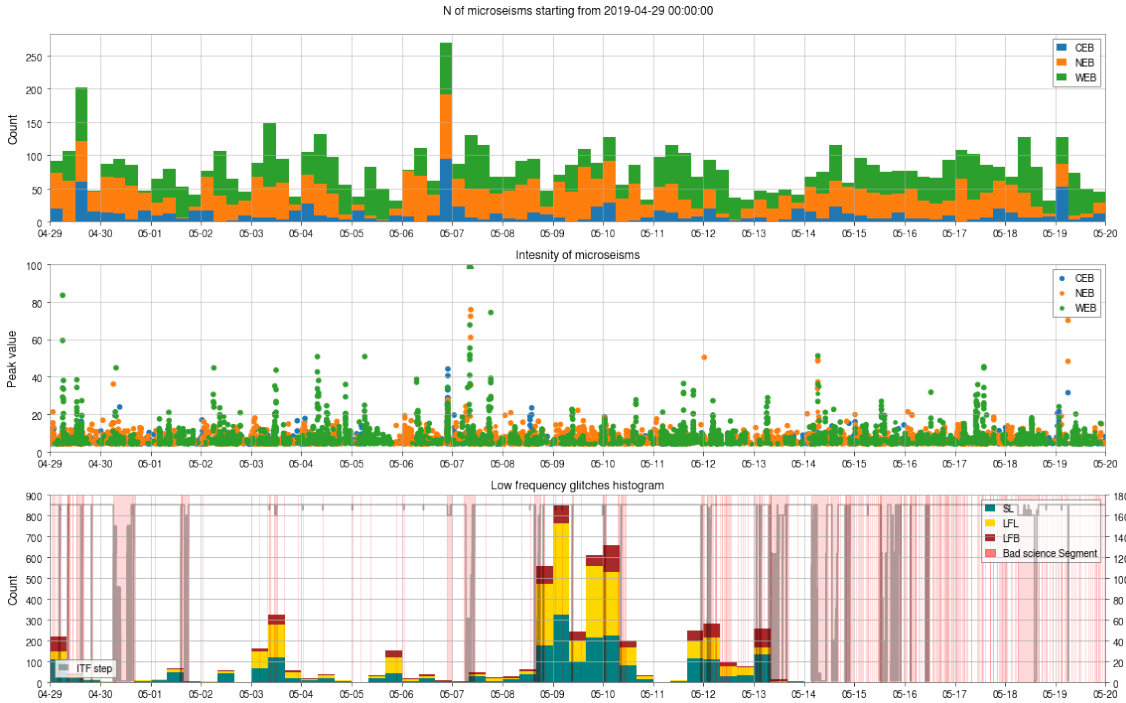

Figure 4.12: Timeline of the seismic triggers during the three weeks concerning the analysis. The top graph is a histogram presenting the distribution in time of these events. The middle graph presents the intensity of these seismic events, as detected by my BLRMS-based algorithm, while on the bottom a plot that showcases the quality of measurement of the interferometer, by plotting the low-frequency glitches distribution in the Gravity Spy dataset and the ITF index. Although no strong correlation between glitches and seismic events is present, spikes in the seismic activity correlate pretty well with some unlocks of the interferometer.

Some summary statistics of the whole dataset are present in fig. 4.13. For all of the buildings the signal in the spectrum seems to be stronger in the vertical direction, both for the seismic noise and for the Superattenuators. Similar to the previously discussed effects, this is probably due to the lower background noise in the vertical direction, evident from the PSDs in fig. 4.2, making 4σ events more frequent, since they need less energy. This effect is probably due to the properties of the soft soil on which Virgo is built, spreading seismic waves not as good in the vertical as the horizontal component. A similar phenomenon is thought to occur in the horizontal plane, where the East-West direction is consistently showing higher SNR than the North-South one, again this trend is roughly maintained for all buildings. When looking at the BLRMS streams that triggered the acquisition of the data, it appears that most of the small earthquakes happen in the higher frequency bands. Activity under 0.5 Hz is more stable, and the triggers of this region seem to be mostly clustered together in time. By correlating these low-frequency swarms with seismic warnings, it appears that these are the direct effect of distant loud earthquakes, that induces long

periods of low-frequency ground motion. These correlations will be better explored in the next chapter.

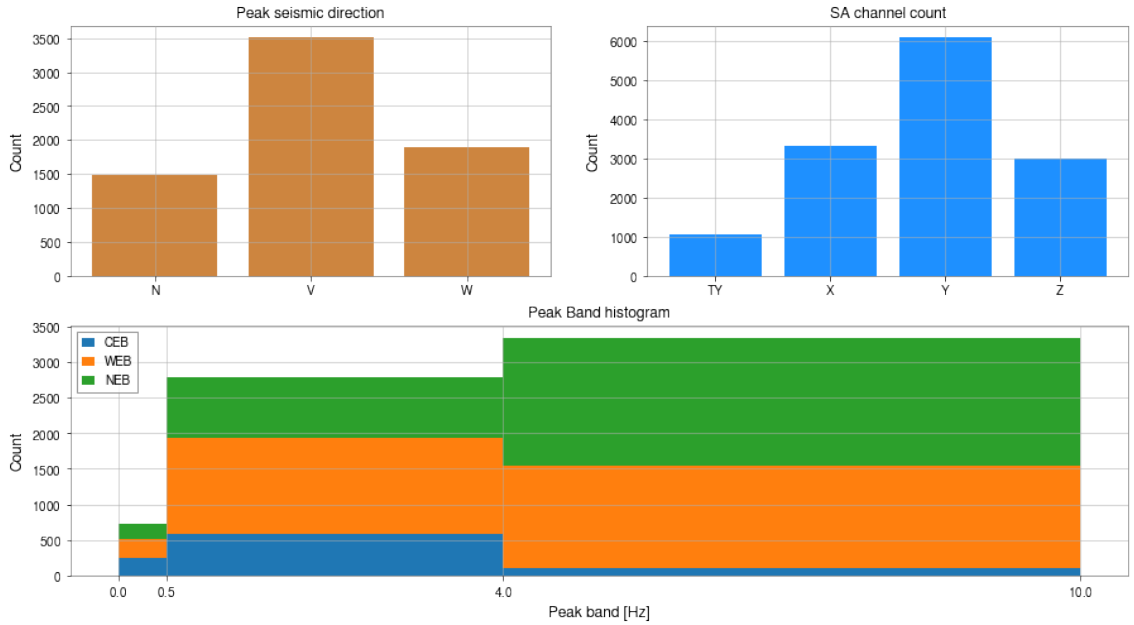

Figure 4.13: Statistichs of the whole dataset, representing the seismic peaks direction, the peaks channels of the F0, and the peak bands of the BLRMS algorithm divided per building.

Interferometer quality

To monitor how the interferometer was performing during these events, and so to eventually check if the Superattenuator anomalous movements had any effect on the measurement quality, a few different parameters were studied. One of the most important things that were monitored is the actual state of the interferometer. Whether the interferometer is unlocked, in the process of locking, or functioning correctly affects how the active feedback loop control keeps the F0 in place. Also, high seismic activity is linked to unlocks and mode changes of how the interferometer operates [45] so this number must always be kept under control. When the interferometer is locked and capable of producing an output stream, a measure of the sensitivity of the instrument is given by the Binary Neutron Star (BNS) range. This represents the average distance in megaparsecs for a standard BNS merger to create a signal in $h(t)$ with an SNR of 8 or greater. This quantity depends on the PSD of the strain channel and it is strongly linked to external noise sources. In fact, this quantity has quite evident daily and weekly fluctuations, due to the effect that human activity has on the detector. A plot of the status of the ITF and the BNS histogram is present in fig. 4.14 and the evolution of the BNS range over a week is present in fig. 4.15.

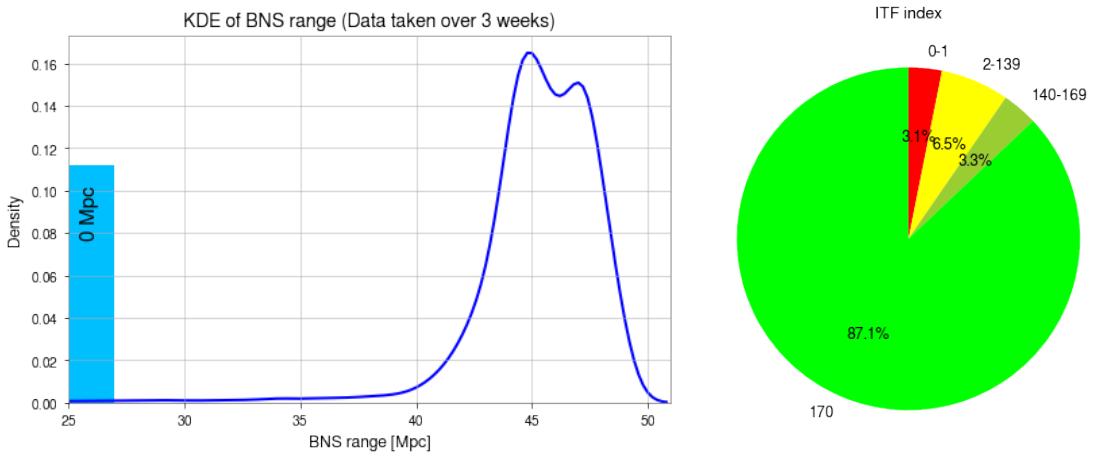

Figure 4.14: On the left, the Kernel Density Estimate (KDE) of the BNS range, on the right, a pie chart of the ITF index channels, both of these were calculated for the 3 weeks of the analysis. The uptime of the Interferometer during this period was just under 90%

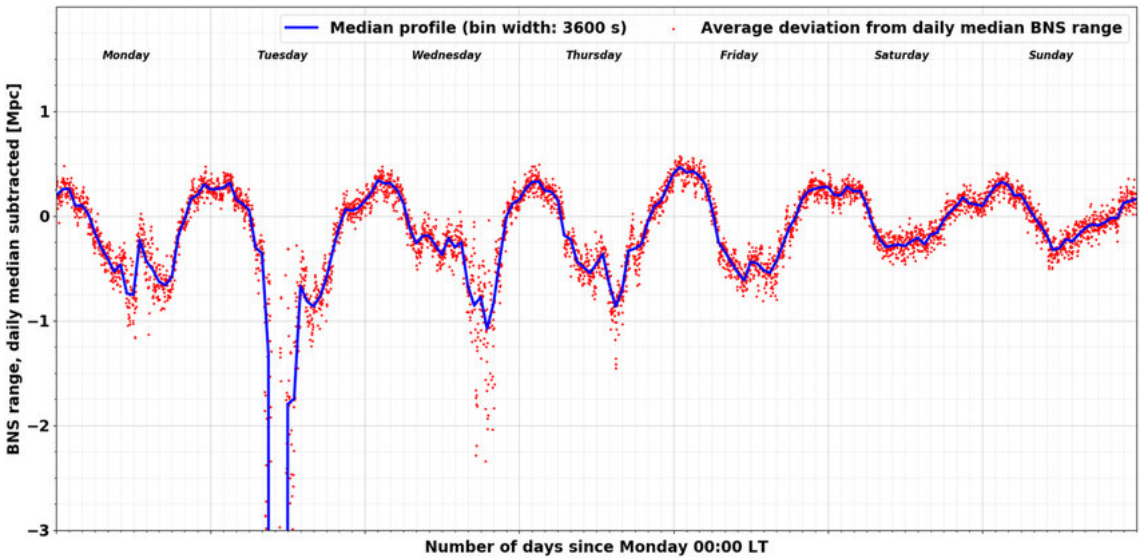

Figure 4.15: Evolution of the mean BNS range as the work week goes on, taken from [45]. It is clear how human activity impacts the detector quality. The drop on Tuesday morning is caused by the weekly maintenance.

Another interesting statistic that needs to be taken into account is if some of the clusters had higher coincides with glitches events. Considering the period of interest the probability that a glitch would happen in a random 16 second window is around 6 %. If some cluster

presents a coincidence value higher than 6%, then there might be some correlations between the Superattenuator movement and the glitches in question.

Wind speed and direction

Wind speed is known to be linked to periods of elevated seismic activity and bad interferometer sensitivity [45]. It was deemed worthwhile to monitor the wind direction during the three weeks of the analysis. A weather station positioned near the CEB monitors the wind speed and the wind direction at all times. In fig. 4.3 there is a summary plot of the wind direction during the 3 weeks of the analysis. One of the goal of the thesis was to use this as a baseline to catch if some correlations between the clusters and high wind activity were present, but no strong winds events were registered during the analysis time span. Some weak correlations with stronger winds were indeed found, but it speculated that these are actually spurious correlations with the day-night cycle. It is a known fact that the wind pattern at Virgo is correlated to a day-night cycle, inverting its direction (towards NE during the day, towards SW during the night) and changing its intensity, reaching higher speeds during the light hours.

4.6 VAE architecture and pretraining

Since the spectrograms are 128x256 pixels for each of the 4 channels, the size of the data to analyze can become quite cumbersome for a clustering algorithm. This is why a VAE architecture was chosen to reduce the dimensionality of the data without losing any relevant information.

The VAE architecture was built using the python machine learning library TensorFlow, through the Keras API [58]. These tools provide an easy and straightforward way to both build and train the machine learning models.

The basic structure chosen for the VAE is a symmetrical encoder-decoder structure, with 5 convolutional layers and a fully connected one at the end. As standard procedure with VAEs, the encoder outputs 2 vectors, one for the μ and one for the $\log \sigma$, which are needed to characterize the Gaussian probability distribution for each sample.

CNNs can process multi-channel input data by constructing different kernels for each channel and then adding them together to construct a single feature map. These usually handle fine multi-channel data like RGB natural images but this has proven to not work properly with the data of this analysis. The 4 channels that represent each sample might be uncorrelated and the network would need to learn the right kernels to keep them separated. During early training the network had trouble distinguishing the different channels, predicting as output the same image

repeated 4 times. Probably with a bigger dataset, longer training time and a bigger network the results would eventually improve, but a much more computationally cheap strategy was implemented instead, based on the power of transfer learning.

Ja-Net

An identical VAE was first trained on the same problem but faced with only one channel at a time, treating the x, y, z, θ_y channels as different samples. The training dataset is now 4 times bigger while the problem complexity is 4 times less. The network in this case came to convergence much faster.

The encoder has been chosen to have 5 convolutional layers, with each layer having double the amount of feature maps as the previous one, going like 4,8,16,16. The exception is the last one, where the doubling was skipped to avoid an explosion in the number of weights of the network. The kernels were of size 8x8 for the first layer and 4x4 for the subsequent ones. To reduce the size of the feature map from one layer to another the stride parameter of the convolution was modified. The first layer used a stride of 4 pixels while the subsequent ones used just 2 pixels, halving each dimension of the feature map at each step of the way. Starting from a single 128x256 image, the output of the last convolutional layer are 16 4x8 feature maps. An example of the feature maps produced by the encoder network at each layer is reported in fig. 4.16.

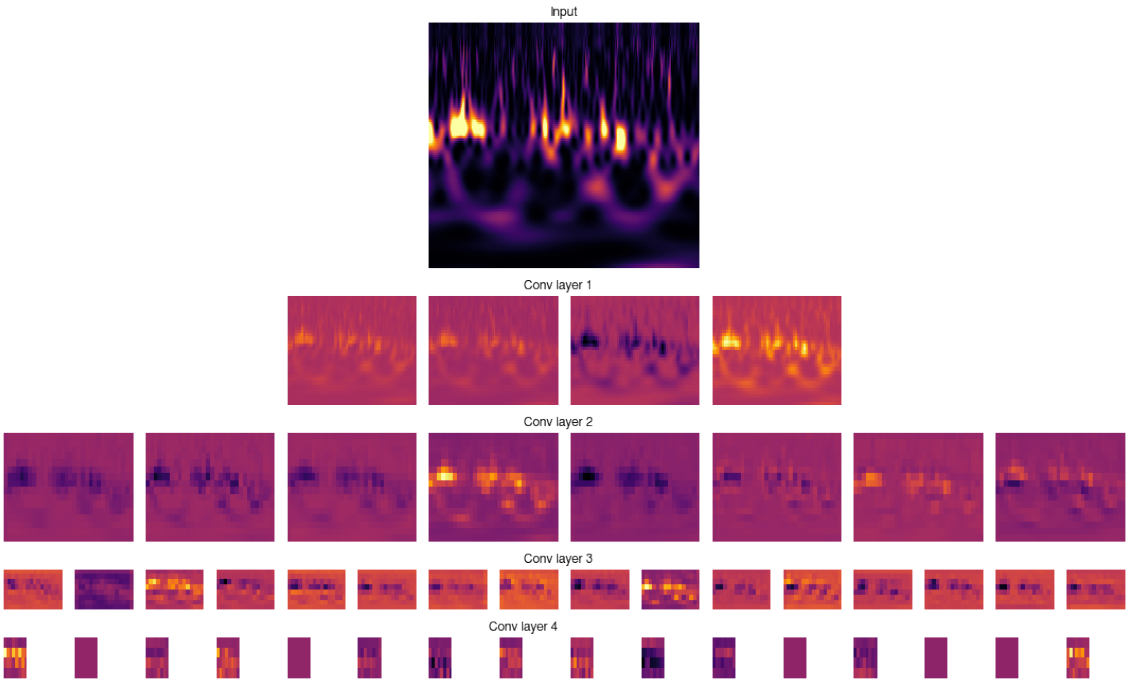

Figure 4.16: Feature maps produced by one of the Ja-Net encoder legs. Convolutions strides make the images smaller at each pass

Once the features maps are obtained, the tensor containing them is flattened and connected to a fully-connected layer with 128 nodes. Finally this layer is connected to two parallel fully-connected layers that will output the μ and the σ of the Gaussian. The network actually outputs the $\log(\sigma)$ value, since this gives the network more fine control in assigning small σ values. Finally, a layer is implemented to take random samples from the probability distribution defined by the last layers of the encoder. A summary of the network is reported in fig. 4.17.
What is obtained at the end is the latent representation of the sample. For this project the latent space was set to have 64 dimensions. The decoder will have a symmetrical structure, where the convolutional layers are just replaced with deconvolutions.

To train the network faster, a technique call batch normalization was implemented. Since at each iteration the mean and variance of the outputs of a given layer can change considerably, the subsequent layers have to learn how to adapt to these changes. This can waste a lot of optimization "power" and batch normalization is one of the options to reduce this cost. Each output is recentered and scaled, so that each mini-batch has 0 mean and unit variance. The subsequent layers will be presented with this standard distribution and will not have to learn to adapt to global changes in the previous one. To not hinder the expression capabilities of the network, an additional rescaling and resizing parameter is learned alongside the network, that

will be used during the inference phase. This technique sped up the learning quite a bit, but it was found that if used at every layer it lead to instabilities in the algorithm, so its present only between a few of them.

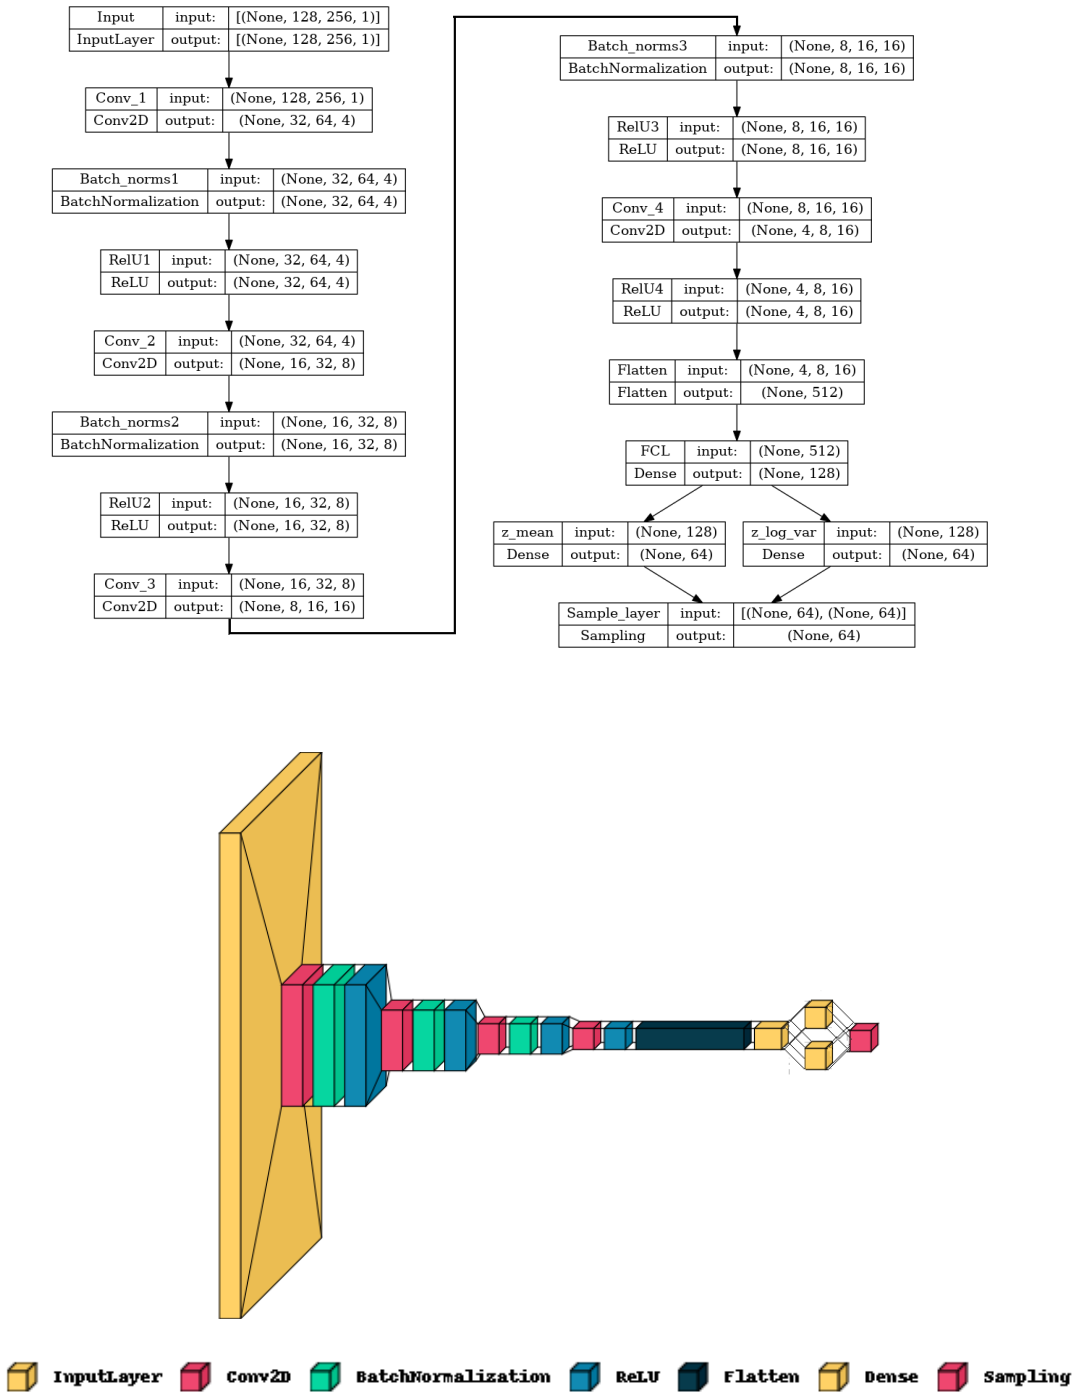

Figure 4.17: Single channel VAE architecture overview. The two parallel layers that are before the sampling are the two that produce the μ and $\log \sigma$ for the samples. No activation was used after the sampling layer, so it outputs directly to the latent space.

Once trained this network will be able to handle single-channel spectrograms. This is not enough for the project, since the objective is to handle 4 channels at the same time, but a valuable resource is gained from this simpler network: the kernels of the convolutional layers. These kernels have been trained to extract spectrogram features so now they can be used for tasks that needs these capabilities, like the more complex network.

The solution that was to found work best for handling the 4 channels is to construct a network by copying the convolutional layers of the single network 4 times in parallel, concatenating their outputs together and linking them to the same fully connected layer, which will translate the information coming from the 4 streams to the same latent space. Like the mono-channel case, the decoder has a symmetric structure to the encoder, making the whole structure of the network roughly resembles the Cyrillic character Ж, which has a sound similar to the English "j". For this reason, it was decided to call the architecture Ja-net. An overview of the whole structure is found in fig. 4.18. This structure has been inspired by the one used in [59]. In this paper, the authors use a particular VAE structure to encode heterogeneous 2-dimensional channels. They tackle the problem by constructing an independent encoder-decoder pair for each channel studied, but all of them talk to the same latent space. For this algorithm, the architecture was simplified, by actually joining together the independent encoder-decoder pairs to a fully connected layer before entering the latent space. This results in a new network that can achieve much better performance with much less computational burden. The possibility of computing parallel information about different channels is what makes the network able to handle uncorrelated channels. This architecture has the advantage that it can be easily scaled up, by adding more and more parallel channels. Some preliminary work was done by adding the accelerometer channels to the architecture, bringing the total number of parallel data-streams upto 8. This showed promising results and will be the focus of some of the works planned for the future.

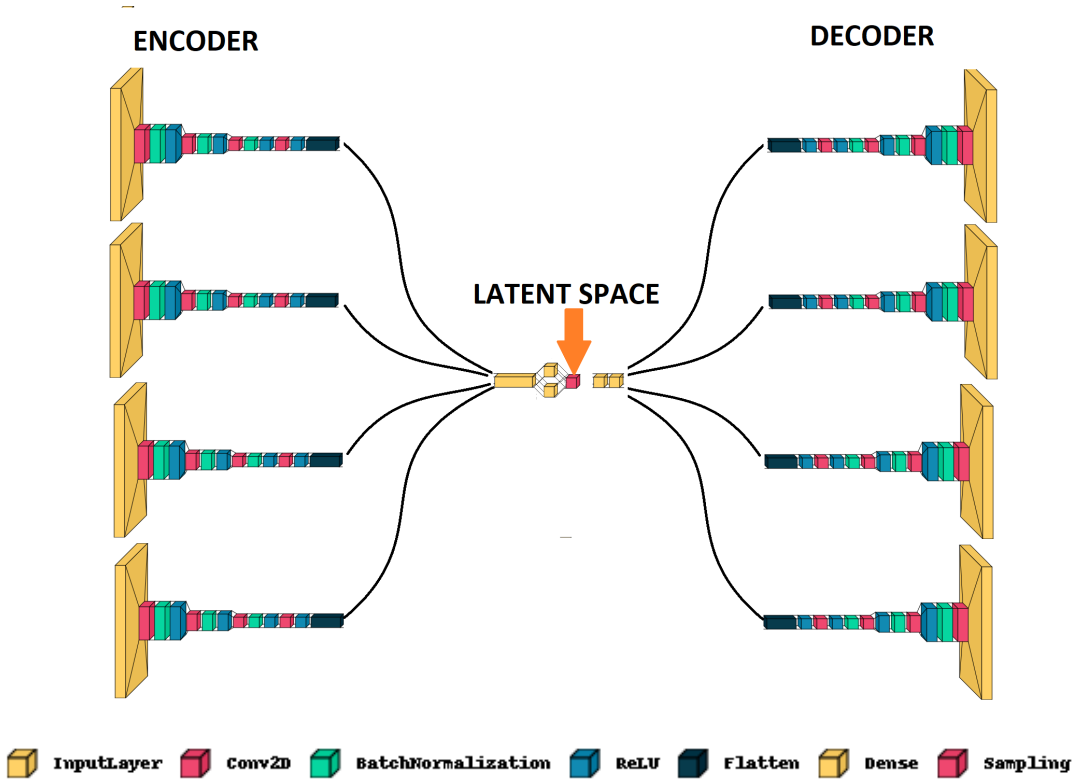

Figure 4.18: Schematics of the multi-channel VAE. Each single "leg" of the network is inizialized with the kernel weights of the mono-channels

The latent space of this VAE has 64 dimensions, like the single channel one. The Kullbach-Liebler divergence has been given a weight equal to 1.5 and the single network has been trained for 20 epochs, while the multi-channel one for 40. Adam was the optimizer of choice with a learning rate of 0.02 and the loss function that measures the distance between the input and the decoded image was the MSE. This function is known to give blurry results as output, but since the Q-transforms already look like undefined blobs this is not much of a problem. The BCE loss function was also tried, with some promising results, but the final total loss appeared too great, making the KL therm much less prominent, which lead to a bad characterization of the latent space. The training data was passed through the network in batches of 32 images each, randomly shuffled at each iteration to reduce biases. The training data was randomly selected and comprised by 80 % of the data, while the remaining 20 % was in the validation set to check if the algorithm showed any overfitting.

All of these parameters have been found by some manual hyperparameter searches. With unsupervised algorithms is really hard to implement some kind of automated hyperparameter

searches since the loss function alone does not cointain any information about the conditioning of the latent space, which is a fundamental quantity to check. Each training of the Multi-channel VAE took roughly 1 hour, while the pre-processing step from Superattenuaotrs time-series to Q-transforms took around 40 minutes.

4.7 Clustering algorithm

After the data points have been compressed to just 64 dimensions, the clustering algorithm can be set up. For the project, the GMM algorithm described in sec. 3.4 was found to yield the best performances. One downside of this framework is that it needs to know in advance how many Clusters it needs to find. The number was arbitrarily set to 40, so 4 classes per Superattenuator. Not many experimentations were done to fine-tune the number of clusters since this would run the risk of "overfitting" over the dataset, so the initial random guess was considered the most unbiased one. In the future, this algorithm can be improved by a better choice of the class number.

The software implementation of GMM was obtained through the Scikit-learn Python package [60].

Summary

This chapter followed the main steps taken during the thesis, and how they were instrumental to achieving the final objective of clustering together Superattenuators F0 motion. The seismic landscape at the Virgo site was first analyzed in detail, which informed the design choices behind the seismic detection algorithm based on the BLRMS. Then how the time-series data was converted to spectrograms that better highlighted the features was described, whit a focus on explaining the principles behind the Q-transform. After the dataset of Superattenuators spectrograms was constructed some analysis of the data was conducted to better understand what possible biases could the algorithm be faced. The complex nature of the problem and the limited computational resources at my disposal required the development of a novel VAE architecture, Ja-Net, to handle 4 spectrograms at the same time for dimensionality reduction. Finally, the specifics of the clustering algorithms used conclude the chapter.

Chapter 5

Results

In this chapter, everything introduced so far will come together, and the results of the work done within the thesis will be explored. Sec. 5.1 is dedicated to assessing the performances of the VAE architecture, by inspecting the algorithm decompression quality and the conditioning of the latent space. In sec. 5.2 the main objective of the thesis is pursued by analyzing the patterns found by the clustering algorithm GMM. The final goal was to determine if these patterns could be linked to some specific origin or to how they could be linked to the detection quality. Sec. 5.2.2 explores if the algorithm is able to cluster together spectrograms coming from the same superattenuator, to gain insights into the noises and the different characteristics that the different suspensions might have. How these patterns are linked to the detector performance was studied in sec. 5.2.3, where links with unlock events, glitches and drops in the BNS range are studied. To test some hypothesis that came up during the analysis, it was decided to conduct some experiments near the suspensions to produce artificial ground motion on the site to trigger the algorithm, and the findings are reported in sec. 5.3. Finally in sec. 5.4 an overview of the possible future applications and improvements of this algorithm is given.

5.1 VAE performances

Before analyzing the results of the clustering algorithm, it is first necessary to assess the performances of the VAE. The specifics of the architecture and the training were described in sec. 4.6, while here only the results are reported. Fig. 5.1 shows the training evolution of the single channel VAE. Both the evolution of the training and the validation data is plotted, as well as the ratio between the KL divergence and the total loss. After the first epoch, the network seems to perform pretty well, minimizing the total loss pretty quickly. The training was stopped as soon as no improvement over the validation set was found. In fig. 5.2 how the decompression performance evolved during training is reported. The reconstruction improves as the network trains and the algorithm learns to encode more and more features.

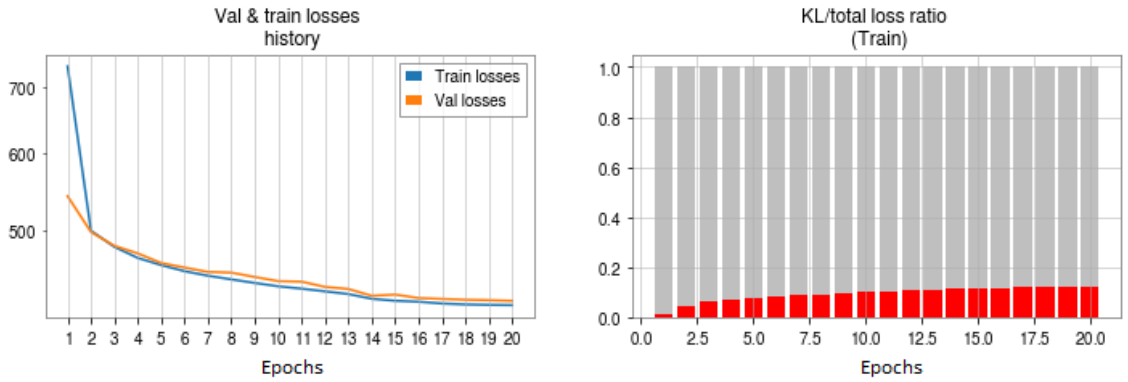

Figure 5.1: Loss evolution of the mono channel VAE. On the left, the evolution of the train and validation losses is present while on the right the ratio between the total and the KL component is reported.

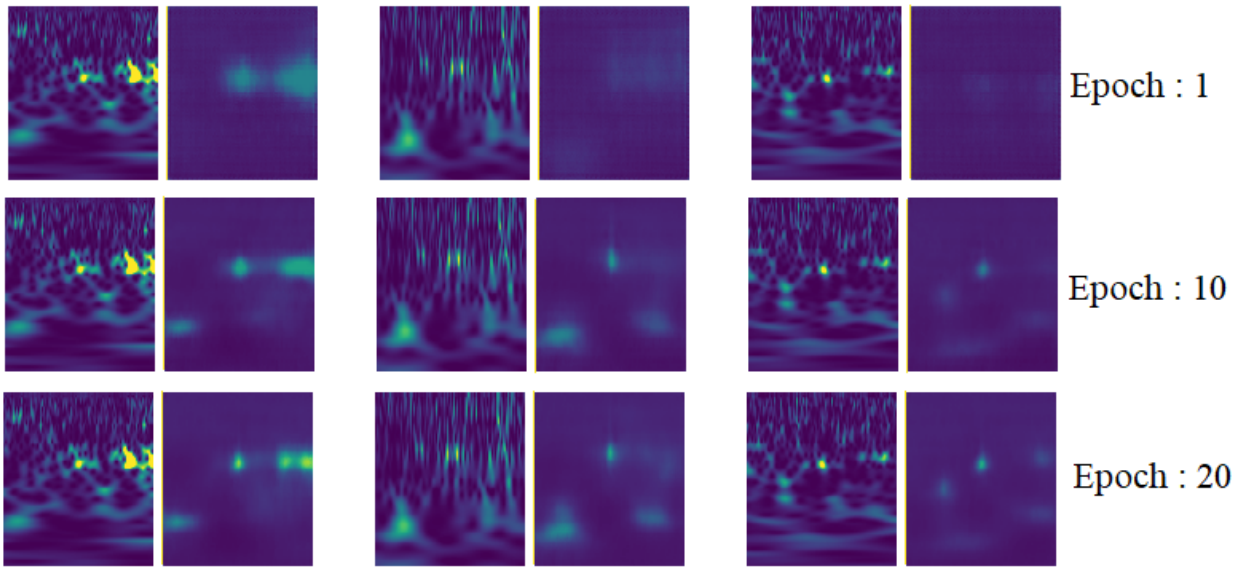

Figure 5.2: How the performances of the single channel VAE evolved during training. On each column a different uncorrelated spectrogram is plotted, on the left, the original data is reported while on the right the decompressed data is present. Data is taken, from top to bottom, from epochs 1, 10 and 20. It is clear how reconstruction quality improves with epochs

It was found that during training the optimizer tried immediately to minimize the KL divergence, reaching levels near ~ 0 even before the first epoch was over. In this step, most of the predicted spectrograms do not look clear and only the vague morphology of the image is being reconstructed. As the training goes on the algorithm learns how to trade between an increase in the KL divergence and a reduction in the reconstruction loss, leading to an overall performance

improvement. This pattern continues until convergence is reached. The train and validation losses appear to be quite similar, so the network is not at risk of overfitting.

After the single-channel VAE, was trained, the kernel weights were used to initialize the multi-channel VAE Ja-Net. The network came to produce much clearer results much faster, as expected due to the transfer learning step. One peculiar behavior is how quickly the new network produced outputs where the spectrograms reached the saturation point, set at 1, as it is pictured in fig. 5.4.Even after continuing the training with the single channel VAE, this effect was not witnessed, while after just 2 epochs of Ja-Net, the network came to saturate the outputs.

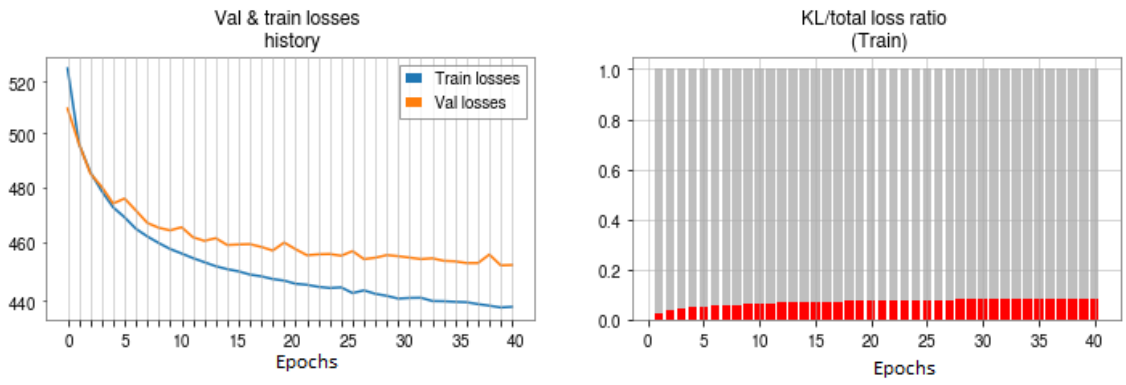

Figure 5.3: Evolution of the training and validation losses for the final multi-channel VAE (left) and the evolution of the ratio between the total loss and the KL divergence component (right).

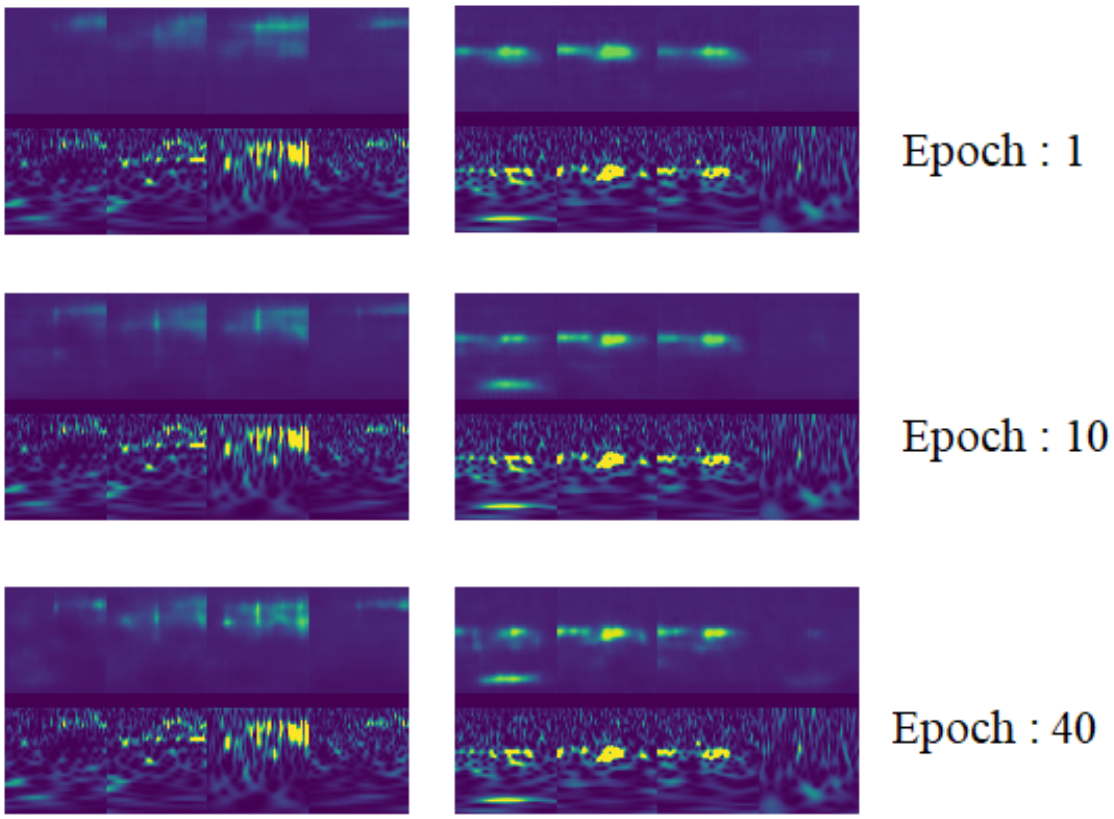

Figure 5.4: How the performances of the multi-channel VAE evolved during training. Each epoch shows on the bottom the original data while on the top the reconstructed version of it. It is clear how the network learns to represent finer and clearer structures as the epochs go on.

Contrary to the mono-VAE, this network seems to be more prone to over-fitting, as it is reported in fig. 5.3. This behavior is suspected to be due to the smaller dataset size when compared to the previous training, which makes it harder for the algorithm to learn generalized features. Still, during the whole training, the validation set showed improvements in the loss function, so the training kept going until no improvements were found in the loss function.

In fig. 5.5 some examples of the reconstructed spectrograms are present. The algorithm manages to preserve all of the main features in the spectrograms while discarding all of the background white noise. Some of the finer features are lost, like quiet low-frequency lines present in the examples, but these behaviors are to be expected since the VAE just learns to recognize the features that appear multiple times in the dataset. If a feature is unique to a sample the VAE will just ignore it.

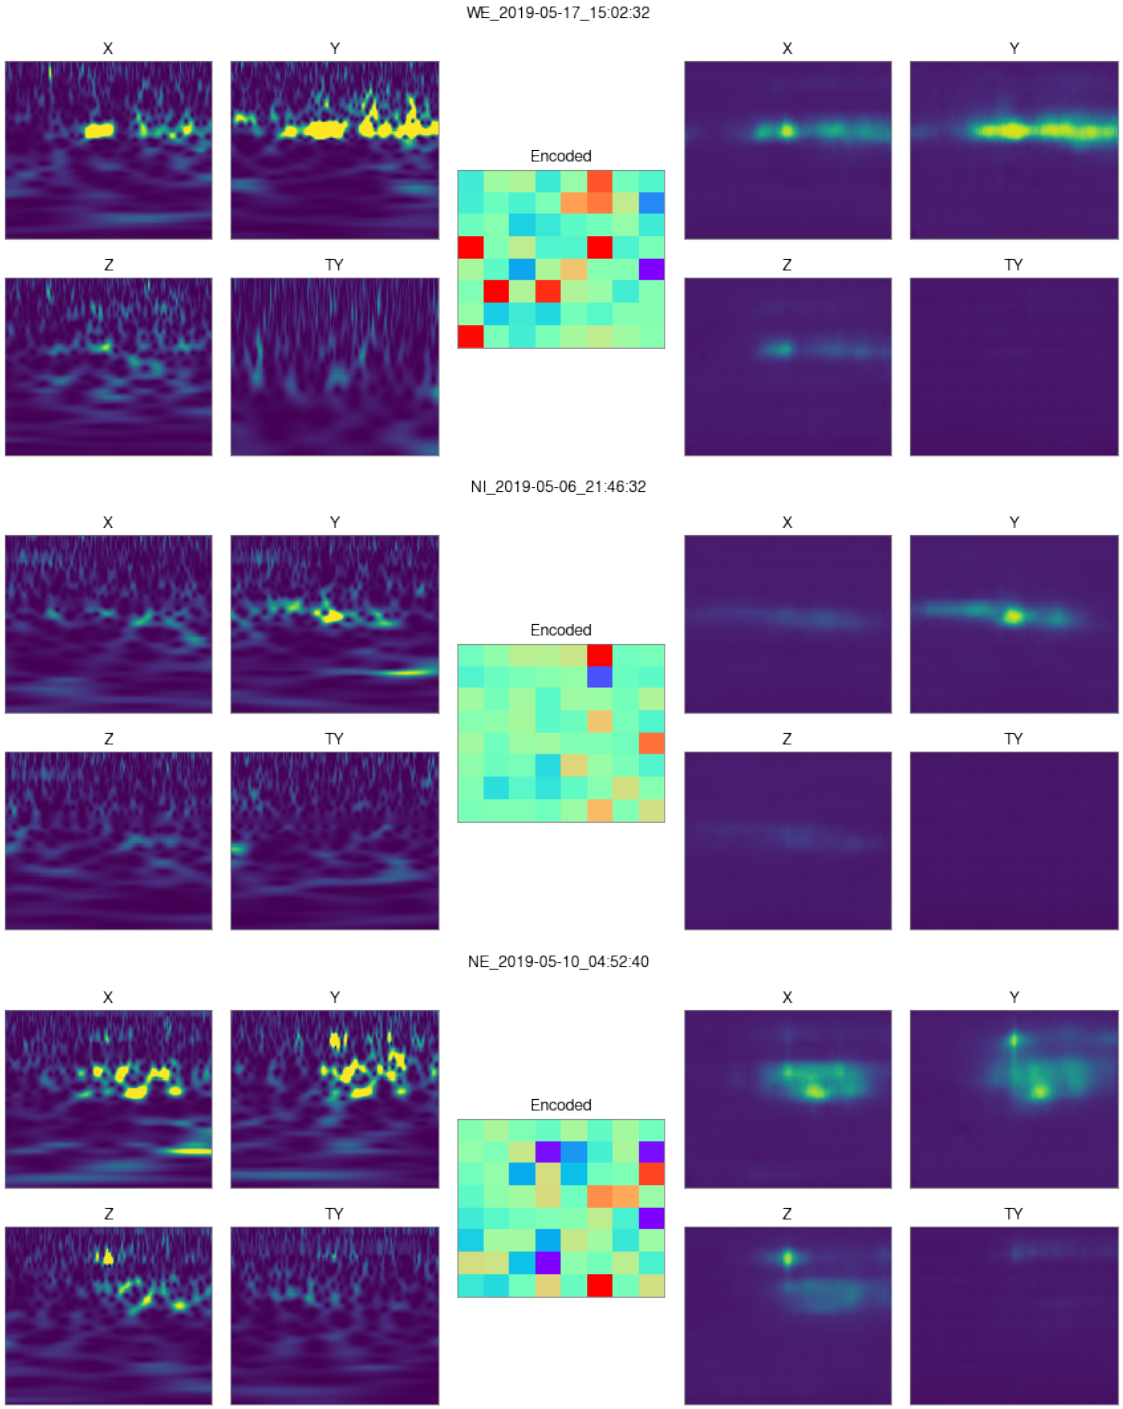

Figure 5.5: 3 examples of spectrograms encoded by Ja-Net. The raw data is on the left, the encoded representation is reported in the center and the decoded spectrograms are on the right. It is evident how the decoded images contain much less information while still preserving the main morphology of the image.

To confirm that the latent space has the desired properties, some tests were conducted. It was decided to encode the whole dataset in the latent space, and then plot a histogram of the values that the latent μ and the latent $\log \sigma$ took.

Normally in VAEs, the features encoded in each specific latent variable are assigned randomly at initialization. To sort the variables in some relevant order it was decided to measure the total influence that a single latent variable has on the final output. The precise calculation of this value is quite hard since the impact of a single latent variable on the output cannot be estimated independently from the others. A good approximation was found by calculating the magnitude of the weights assigned to each variable in the first layer of the decoder. To find the magnitude of the i-th latent dimension the algorithm calculated $\mathrm{mag}(i) = \sum_{j=1}^{K} |W_{ji}^{in}|^2$ where W^{in} is the input matrix of the decoder and K is the size of the first hidden layer. By ordering the latent variables histograms with this criterion, the final plot is presented in fig. 5.6. Variables with the smallest magnitude (upper left part of the figure) have μ close to 0 and σ close to 1: their probability distribution just collapsed to the prior. Since the encoder will output almost the same probability distribution independently of the input, these variables basically do not hold any relevant information. This finding is compatible with the low-magnitude measurement, since their value is suppressed by the input decoder weight matrix.

As the magnitude increases, the μs become more spread out and the σs smaller, so now these variables can encode different features for different samples, actually participating in the reconstruction process.

The fact that during training some latent variables of the VAE just become "inactive" is a pretty well-known phenomenon commonly referred to as over-pruning [61]. The optimizer prioritizes the minimization of the KL divergence regularization therm over the reconstruction loss by making some distributions just collapse to the prior. Even by extending the training, the algorithm is not able to assign any reconstructed feature to these variables. A lot of modifications of the VAE algorithm have been proposed to tackle over-pruning [62] [63], but they were not implemented. During future works exploiting more advanced VAEs architectures will be one of the main objectives.

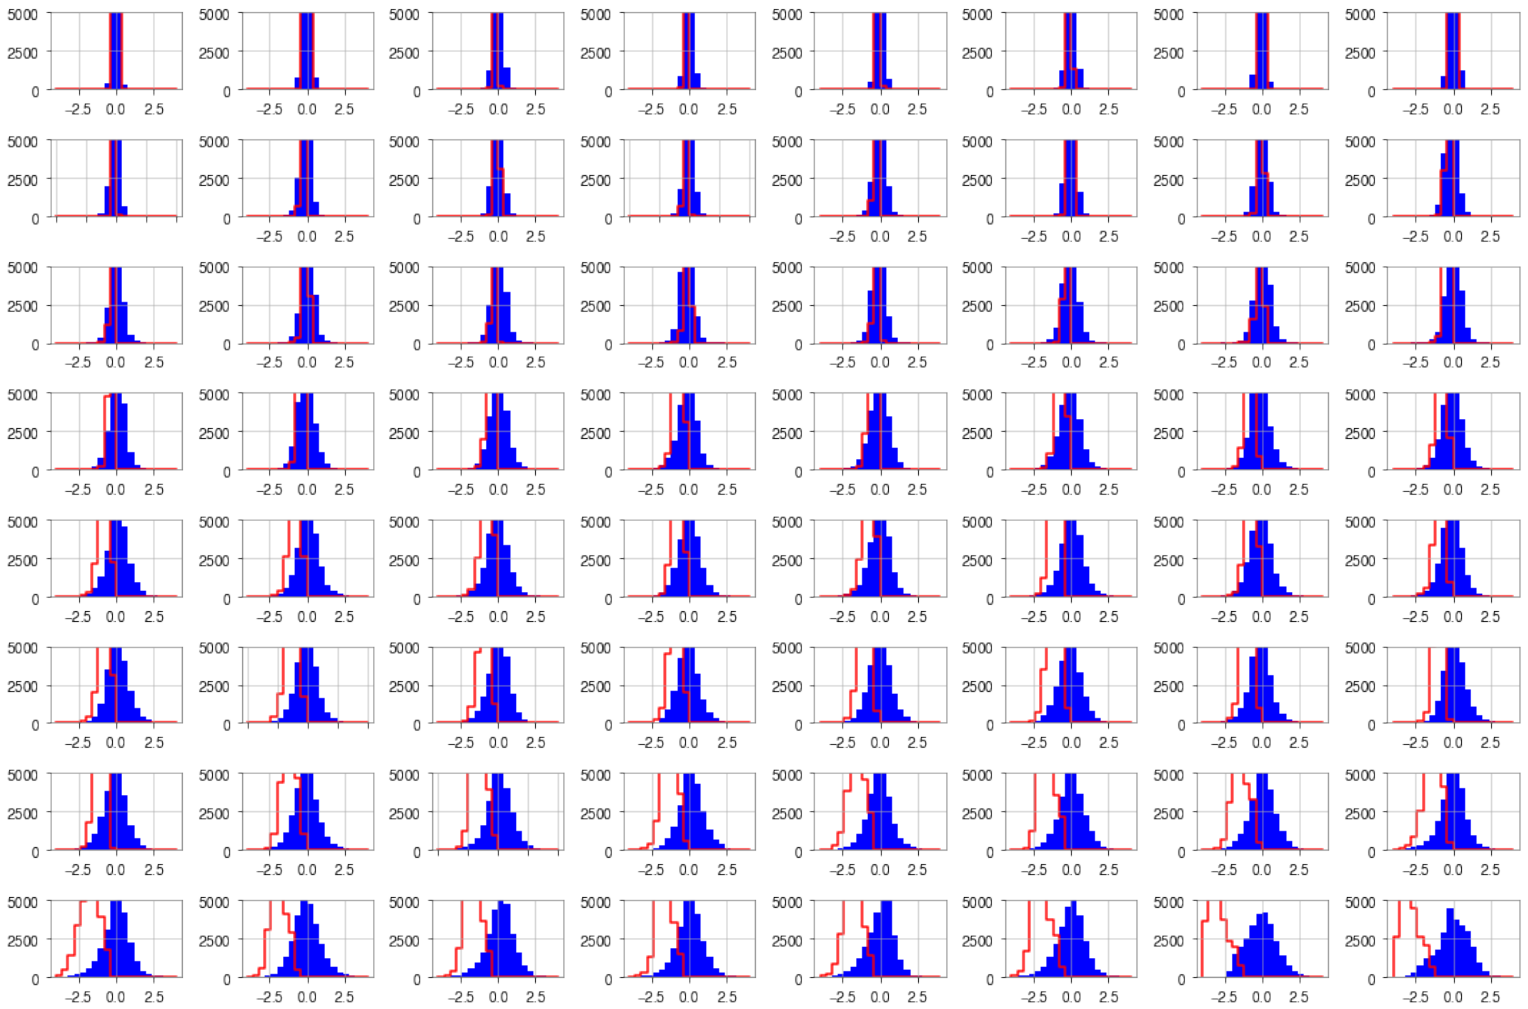

Figure 5.6: Histogram of the distribution of each latent variable for the encoded dataset. The latent variables are ordereb by their $\text{mag}(i) = \sum_{j=1}^{K} |W_{ji}^{in}|^2$ value. In blue the μ are plotted while red represents $\log(\sigma)$

5.2 Clustering algorithm performances

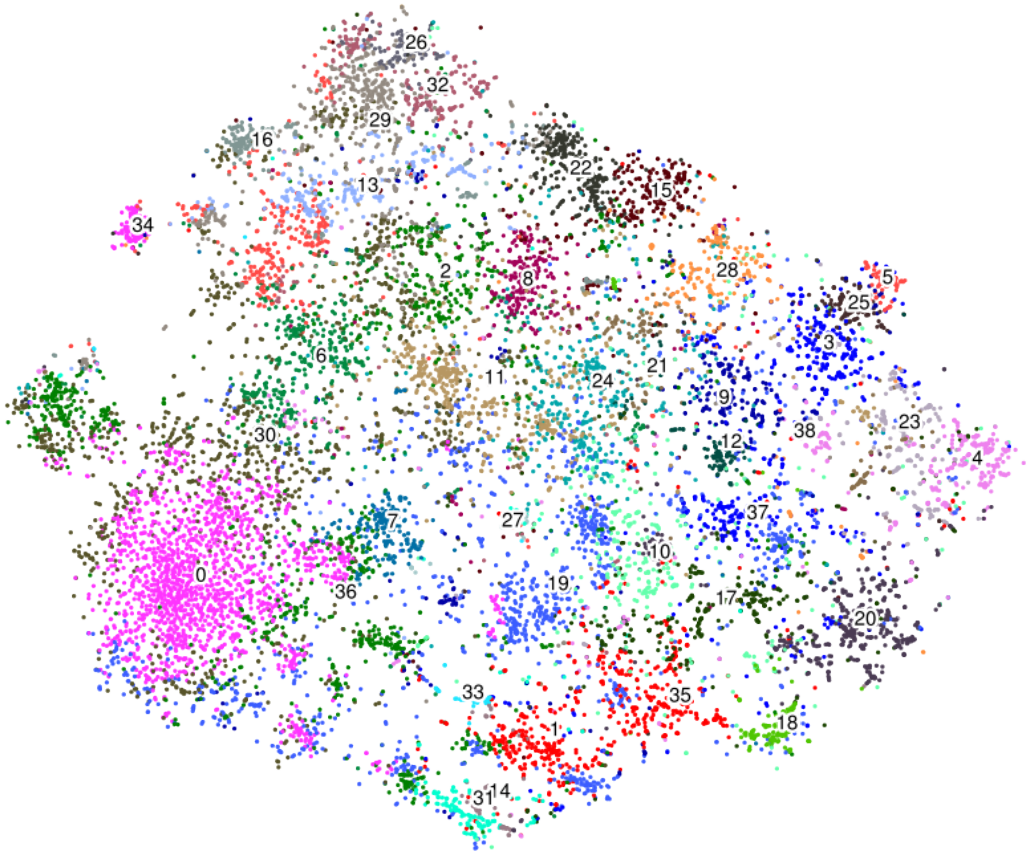

Figure 5.7: The whole dataset projected in just 2 dimensions with tSNE algorithm

This section is dedicated to the analysis of the clusters found by the GMM algorithm. To assess the correlations that these patterns might have with outside factors, the analysis is focused on variables that the clustering and the VAE have no access to, like the original Superattenuator that originated the spectrogram or the quality of the interferometer corresponding to the event or the time of day. If clusters present deviations from the expected value, calculated over the whole dataset, of some of the variables, then a correlation between the patterns and the outside factor can be hypothesized. These deviations from the mean are checked both with thorough Bayesian Null hypothesis tests and the visual inspection of the plots. The Null hypothesis, which serves as a benchmark to test the findings of this analysis, is that the clustering algorithm constructs each cluster from randomly picked samples from the whole dataset. On the plots presented in the rest of this work, the Null hypothesis is always plotted as a red line over imposed with the

actual measurements done for each cluster. The black marks on top refer to the credible interval.

5.2.1 t-SNE

To have a better understanding of how the whole dataset is represented in the latent space of the VAE, I have decided to project it into just 2 dimensions with the tSNE algorithm. The results are presented in fig. 5.7. This algorithm has the task to maintain local relations between points, so the local neighborhood of each point in the encoded original 64-dimensional space should be preserved, meaning that nearby spectrograms in the tSNE representation should look similar.

t-Stochastic Neighbor Embedding (t-SNE) is a widely used method for dimensionality reduction, developed by L.J.P. van der Maaten and G.E. Hinton [64]. It is a non-linear and non-parametric projection method that has been used extensively to visualize highly dimensional datasets in just two or three dimensions, due to its ability to preserve features and long-distance relations even after the projection. The main idea behind t-SNE is to assign a probability distribution to the neighborhood of each point x_i

$$p_{i|j} = \frac{\exp(-\|x_i - x_j\|^2/2\sigma_i^2)}{\sum_{k\neq i}\exp(-\|x_i - x_k\|^2/2\sigma_i^2)} \tag{5.1}$$

This represents the probability that a data point j is the neighbor of i. The σ_i parameters are assigned so that the local entropy $H(p_i) = \sum_j p_{j|i} log(p_{j|i})$ of each data point is the same, and is set to a user-determined number Σ called the perplexity. This is done to assure that the σs will roughly inversely scale with the local density of the dataset so that each point has a similar number of neighbors inside the Gaussian that it casts. This parameter is smaller for higher-density regions of the data space and vice versa. The points are then randomly projected into the lower dimensional space at coordinates $\mathbf{y}_i$, where the algorithm will calculate another probability distribution similar to the previous one:

$$q_{i|j} = \frac{(1 + \|y_i - y_j\|^2)^-1}{\sum_{k\neq i}(1 + \|y_i - y_k\|^2)^-1} \tag{5.2}$$

This has a much longer tail, which has the property of preserving short-distance information while pushing away points that are far apart in the original space. Finally, the objective of the algorithm will be to make these two probability distributions. This is achieved by minimizing the KL divergence between the two distributions, with an optimization procedure. Finally, a low-dimensional representation of the original dataset that preserves local relations is obtained.

For the parameter tuning of this implementation of t-SNE a few experiments were run. The perplexity was chosen to be 50, while the inizialization was done with PCA. The metric that t-SNE used to calculate distances between samples was the cosine distance.

Contrary to many other tSNE visualizations, this seems to be less "spotty", and has a more uniform distribution over the whole latent space. This is a direct and expected result coming from the regularization parameter of the VAE, since every probability distribution is "pushed" towards a standard normal, the latent space will tend to be uniformly and densely packed with samples coming from the training set. This is a hint that the VAE is working as intended.

The colors in 5.7 are given to each sample by the GMM clustering algorithm, which was tasked to find 40 different clusters in the Latent space. The actual numbers are assigned arbitrarily by GMM. Points assigned to the same cluster mostly sit near each other in the tSNE representation, which is a good qualitative indicator that the clustering algorithm is working as intended.

The rest of this section will be dedicated to an overview of the most interesting clusters found, with a summary of the useful insights that these give us on the inner working of the Superattenuator.

5.2.2 Unsupervised Superattenuator Classification

One of the goals of the project was to find patterns in the data that are typical of specific Superattenuators: if some clusters contain mostly spectrograms coming from the same suspension, the inspection of its contents would highlight some features of the noise that are unique to that Superattenuator. I found that the algorithm tends to produce clusters containing samples coming mostly from either the WE or the NE Superattenuators. This result is to be expected since these Superattenuators are positioned the furthest away from any other one, making their surrounding environment the most unique.

West End Clusters

The clusters that were assigned the numbers 26, 32 and 16 have the peculiarity of being composed of mostly events coming from the West End (WE), with the most extreme case happening with cluster 26 being composed of only WE spectrograms. This is the only occurrence of a cluster being composed of a single Superattenuator. In fig. 5.10 an overview of these clusters is present, with the position in the t-SNE visualization highlighted on top and the median spectrogram for the cluster plotted in the middle. In fig. 5.8 an overview of the statistichs of one of these clusters is present.

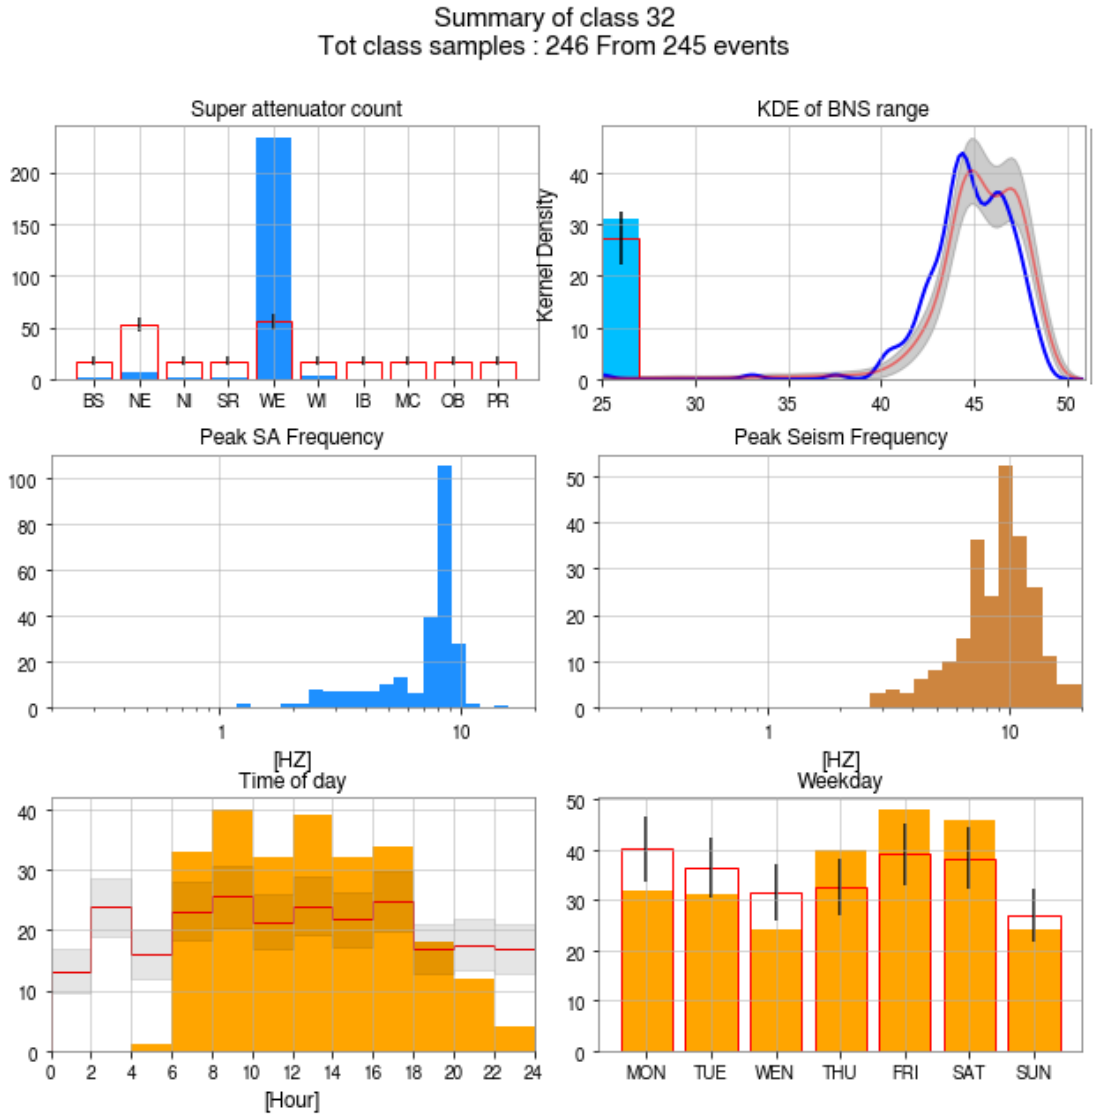

Figure 5.8: Description of cluster 32. This cluster is almost exclusively comprised of WE Superattenuators. The typical frequency at 8 Hz is visible in the histogram above. The events of this cluster seems to be linked with human activity, since they present an evident daily modulation. Also the medium BNS range seems to be lower than the Null Hypothesis, so this cluster might have some links with the bad quality of the detector.

The main features of the cluster seem to be the loud noises at the 8-11 Hz frequency lines, sometimes accompanied by lower frequency features, but with a broader bandwidth and not as loud. It is speculated that the high-frequency lines are originating from some kind of resonance in the seismic pre-isolation stage. In particular, when looking at the peaks in the frequency

distribution of the seismic activity, these seem to be more spread out than in the Superattenuators, like in the middle plots in fig. 5.8. Another compelling piece of evidence comes from the comparison between the median spectrograms for the seismic excitation and the median F0 response, as it is possible to see in fig. 5.9 for the vertical channel for cluster 16 events. The excitation at 11Hz seems to linger on for $\sim 3-5$ seconds after the seismic excitation has stopped. Another interesting note is that these excitations seem to follow a day-night cycle, which could hint at the anthropic origin of the noise, but there is no strong weekly modulation, but a mild dip on Sundays. The cluster is also weakly linked to a lower BNS range than expected, so future works might want to focus on this kind of excitations to better asses their effect on detection.

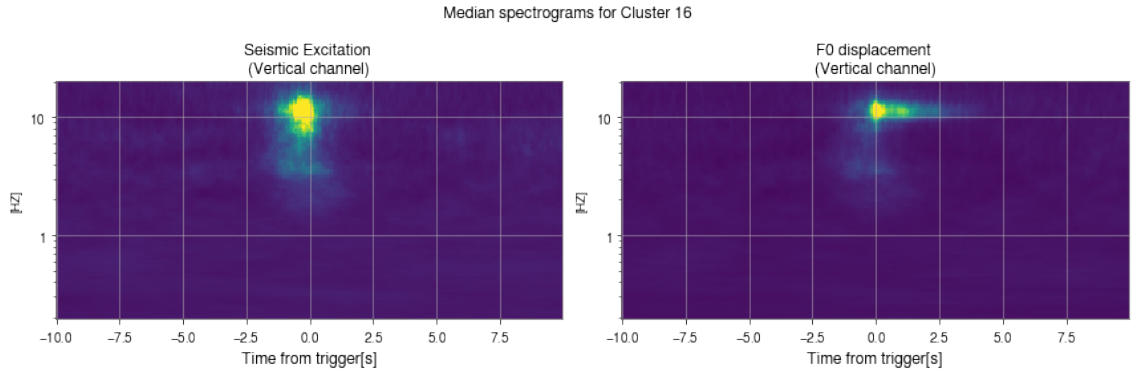

Figure 5.9: Median spectrograms for the seismic excitation (left) and the Filter 0 response (right) for events belonging to cluster 16. The line at 11 Hz seems to linger on in the attenuation system even after the seismic event has stopped.

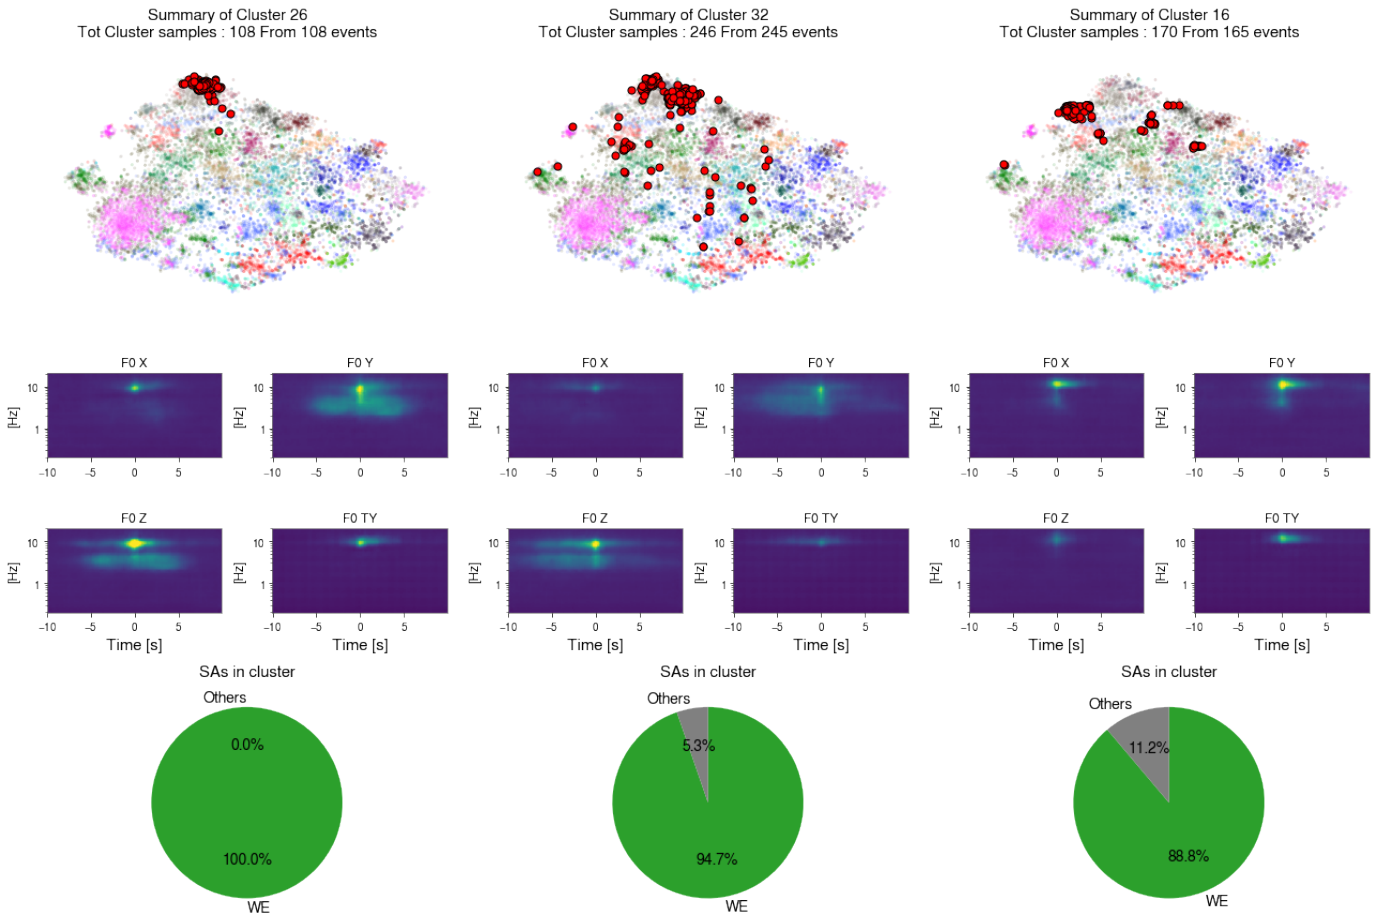

Figure 5.10: Examples of clusters where the WE Superattenuator is the most prevalent one. On top, the position in the tSNE projection is reported, while underneath the median spectrograms for the cluster are plotted. The most peculiar features seem to be these lines between 8-11 Hz

North End Clusters

Similarly to what was discussed in the previous section, there are clusters that contain events coming mostly from the NE Superattenuator. The most prominent ones are shown in fig. 5.12, while in fig. 5.11 a deeper dive into the statistics surrounding cluster 22 is present

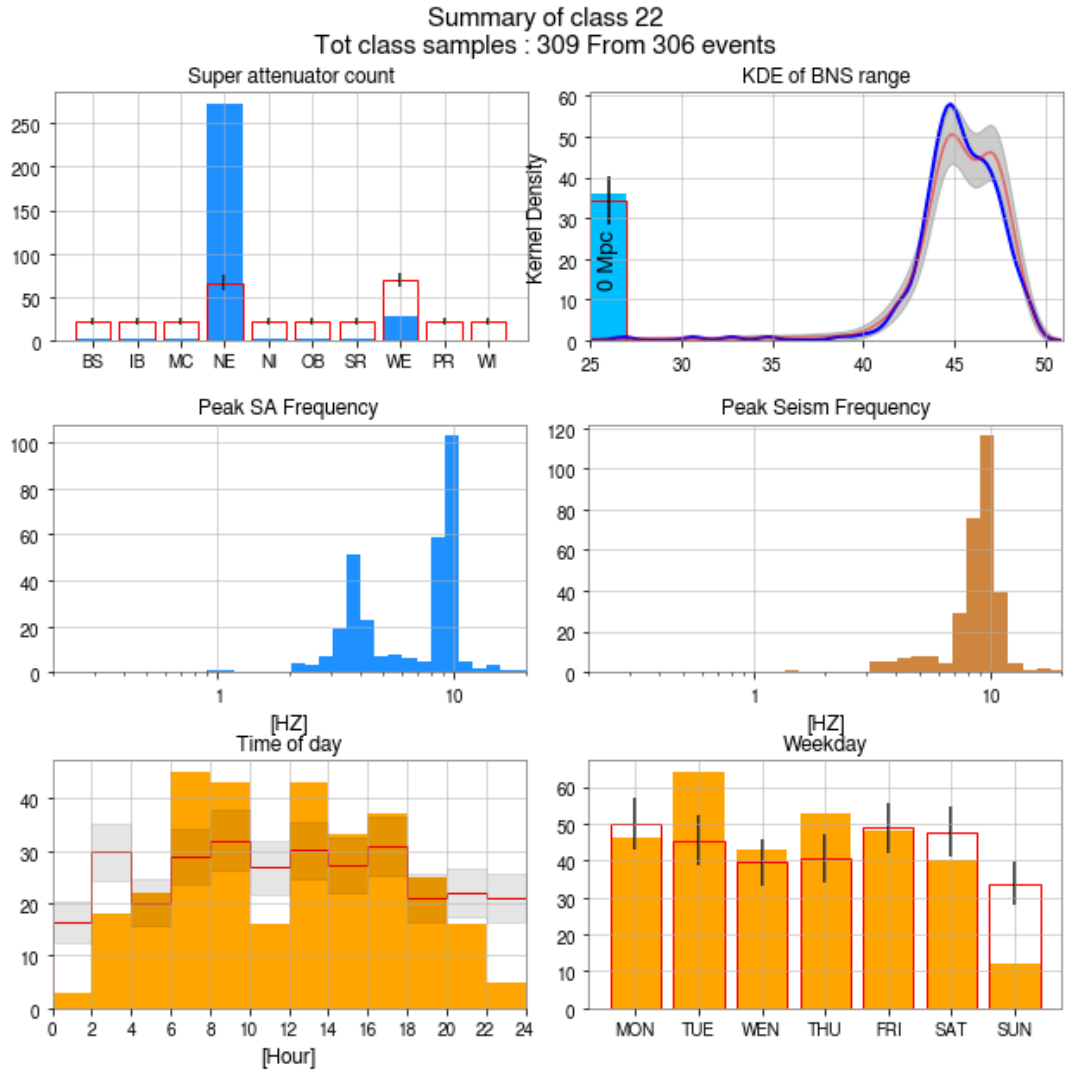

Figure 5.11: Deep dive into one of the NE clusters. The most peculiar features seem to be the doubly peaked frequency distribution at 3.5 and 9 Hz and the daily and weekly modulations that seems to follow the typical working hours. There does not seem to be a correlation with bad states of the interferometer.

The main feature that these spectrograms seem to share is this doubly peaked frequency distribution, at 10 and 3.5 Hz. This frequency distribution seems to roughly be the same in the

3 clusters. By looking at the spectrograms in more detail these clusters seem to represent all similar events but with a slightly different time evolution. Keeping in mind that the spectrogram time window is centered around the peak value of the Q-transform, the algorithm divides these events such that one of the clusters contains spectrograms where the excitation continues after the peak value (cluster 22), a cluster where the peak is at the center (cluster 15) and a cluster where most of the activity happens before it (cluster 8). A GMM algorithm trained with fewer clusters might have grouped these 3 together, since they also appear close in the t-SNE representation.

The time of day distribution of the events in these clusters might give a hint to their origin. They all seem to exhibit both a weekly and a daily modulation, compatible with the work week and the work day: there is a dip in cluster triggers each Sunday and also on the first of May holiday. Previous studies on how anthropic noise affects the seismic landscape at Virgo, has shown that these frequency bursts seem to be linked mostly with seismic excitation coming from the nearby bridges of the FI-PI-LI highway [55]. Normal modes of the bridge at 3 Hz get propagated through the ground to the Virgo Building. Once arrived at the site this seismic motion seems to excite some kind of resonance in the building foundations, making the distinct line at 10 Hz [54]. Even if some further validation work could be done in the future, this seems to be the most probable cause for this cluster.

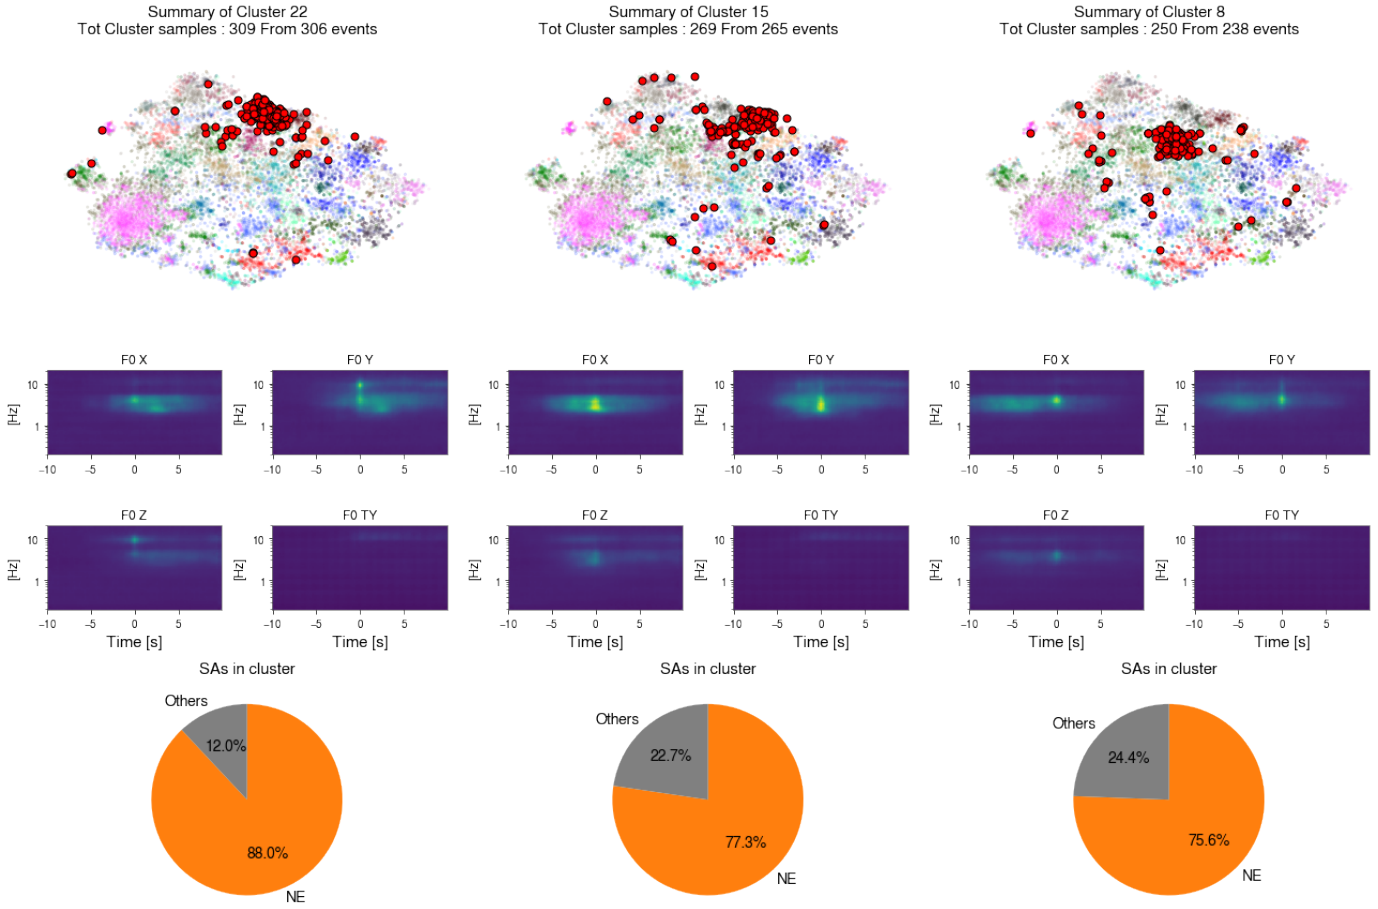

Figure 5.12: Examples of clusters where the NE Superattenuators is the most prevalent one. The most peculiar features seem to be the double frequency lines centered around 3.5 and 10 Hz

Power Recycling Clusters

Usually, spectrograms coming from the Central Building tend to be clustered together, but the Power Recycling mirror seems to be an exception. Cluster 37 is the most prominent example of this phenomenon since half of the spectrograms belong to the PR Superattenuator. Fig. 5.13 presents a comparison between two similar clusters. One is cluster 37, which contains mostly events from the PR Superattenuator, and the other is cluster 3, where the PR is almost absent.

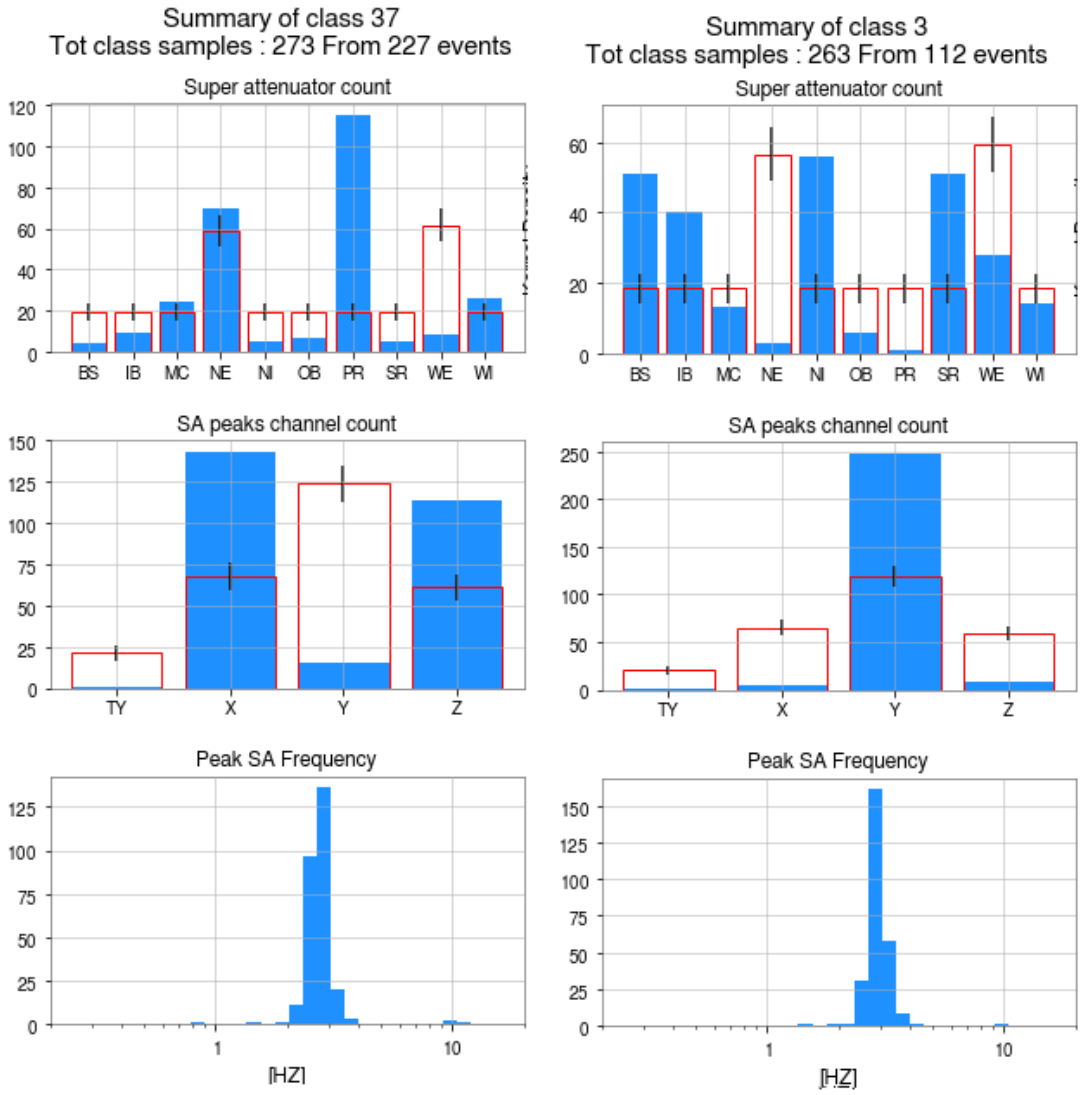

Figure 5.13: Comparison between a cluster that contains mostly events coming from the PR mirror and one where these events are mostly absent. The main difference seems to be the difference in the loudness of the vertical component, almost completely quiet in the PR cluster

The most obvious disparity seems to be the difference in the channel loudness of the vertical component since in the PR cluster is much quieter with respect to its counterpart. By analyzing the raw data in a bit more detail it appears that the noise floor in the vertical control of the PR mirror is much higher when compared to the other Superattenuators, as it is possible to see in the PSD plotted in fig. 5.14. Events that have the same intensity will be more likely drowned in the background noise in the PR spectrograms, while the other Superattenuatorss will exhibit a clearer trace.

After talking with the Superattenuator team at Virgo, the noisiness in the vertical movement of the PR F0 seems to be an already known fact since this mirror needs to be moved much more in the vertical direction with respect to the others. The actuators of the PR F0 will have a larger dynamic range and this introduces more electrical noise. The fact that the algorithm was able to pick up this fact on its own shows promising results for the use of the framework developed during this thesis for further investigation of the noise of the suspensions.

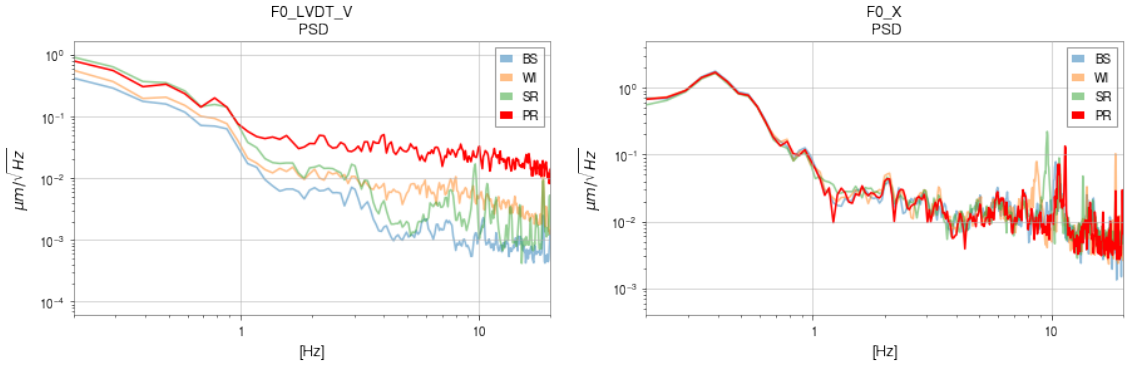

Figure 5.14: PSD of the Power Recycling Superattenuator F0 compared to other Superattenuators present in the central building. for the vertical channel on the left and for the X channel on the right. The elevated amount of noise of the vertical F0 PR channel is pretty clear.

5.2.3 Clusters correlated to the unlock of the interferometer

Periods of elevated ground motion can greatly impact the quality of the detector. To maintain the Fabry Perot Cavities in the resonant condition, actuators put forces on the mirrors to move them back to the dark fringe. During events of higher than normal ground motion, the actuators need to put more work to keep the mirrors stable, introducing great amounts of noise in the detector. If the residual ground motion at the mirror level requires the actuators to operate outside their dynamic range, the interferometer totally loses the lock condition, making the detection impossible. The process of getting back to science mode is quite a time-consuming one [21] and since multiple detectors need to be operational at the same time for gravitational

waves detections to be validated, maintaining the lock status and having a high percentage of up-time is a high priority.

Since ground motion couples to mirror movements mainly through the Superattenuators, I considered it a worthwhile endeavor to explore whether some of the clusters are linked to unlocking events, or more in general to noise in $h(t)$. To get a more precise picture, the ITF index channel, discussed in sec. 2.6, will be used to infer the status of the interferometer at the time of the events. The three clusters summarized in fig. 5.15 are the ones with the strongest link to the interferometer being in an unlocked state. These clusters are luckily quite small and represent some extraordinary and loud events.

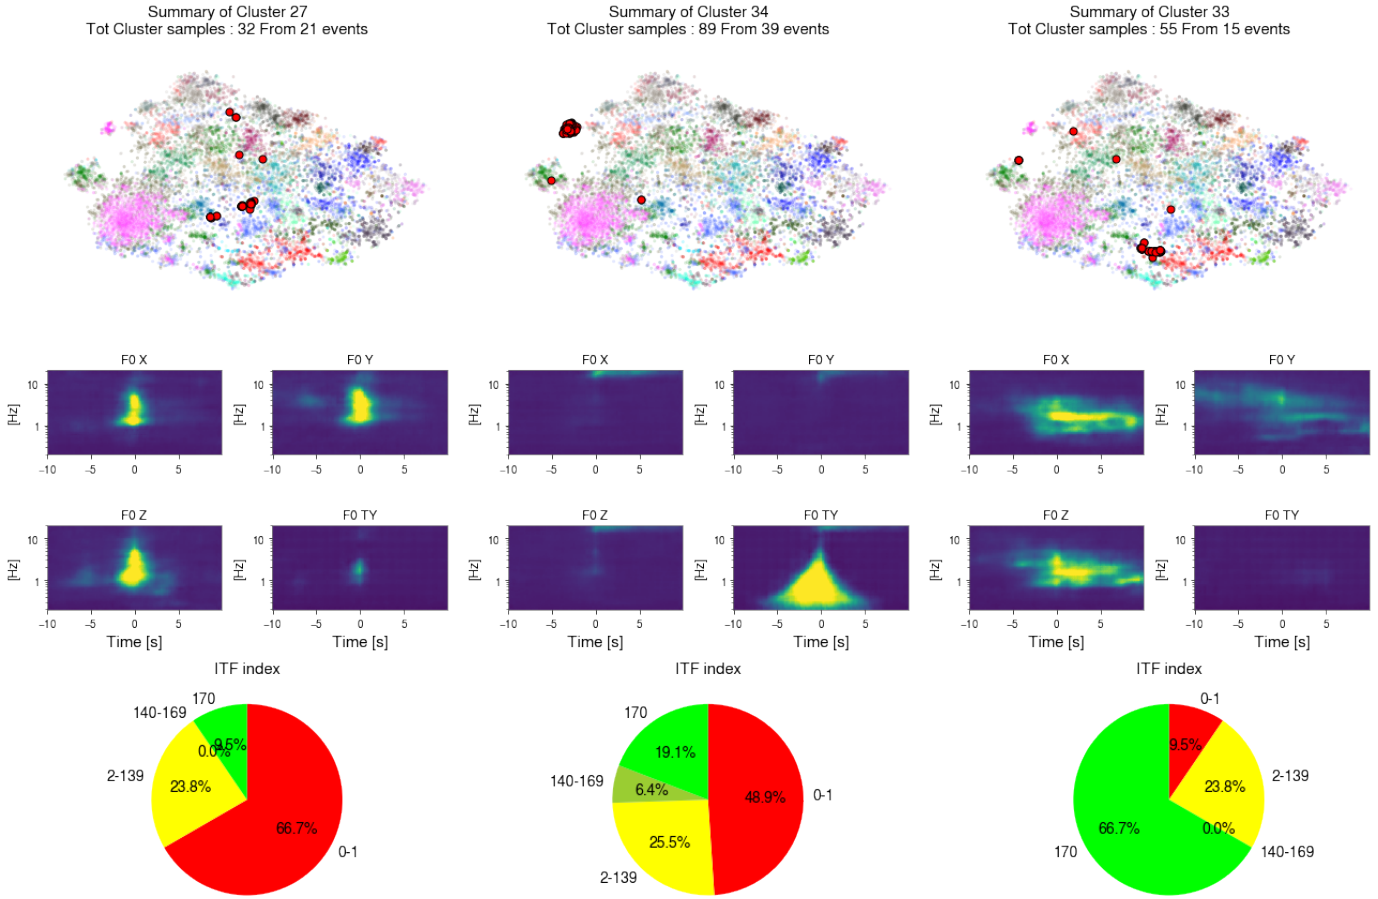

Figure 5.15: The 3 clusters that were linked to the worst performance of the ITF. All 3 of them seem to present some kind of loud events that impact high-frequencies.

Witch hat cluster

The most peculiar of these clusters is cluster 34, which presents a loud signal in the θ_y degree of freedom, that kind of looks like a witch hat, followed by a high-frequency "ring" in every channel, as it possible to see in fig. 5.16. By looking at time series data from one of these events in fig. 5.17, it appears that there is a discontinuity in the θ_y control happening at the same time in all of the Superattenuators. There is also a clear link to the total loss of control of the interferometer. This peculiar behavior seems to originate from the fact that the interferometer output is used in a feedback loop to control the position of the suspension point [25]: whenever an unlocking event happens, the control capability of the suspension inevitably changes, so a discontinuity in the position of the mirror is to be expected. The events that build this cluster seem to be linked to the arrival of distant seismic waves (SEISMON warnings) and activity near the building happening Tuesday mornings.

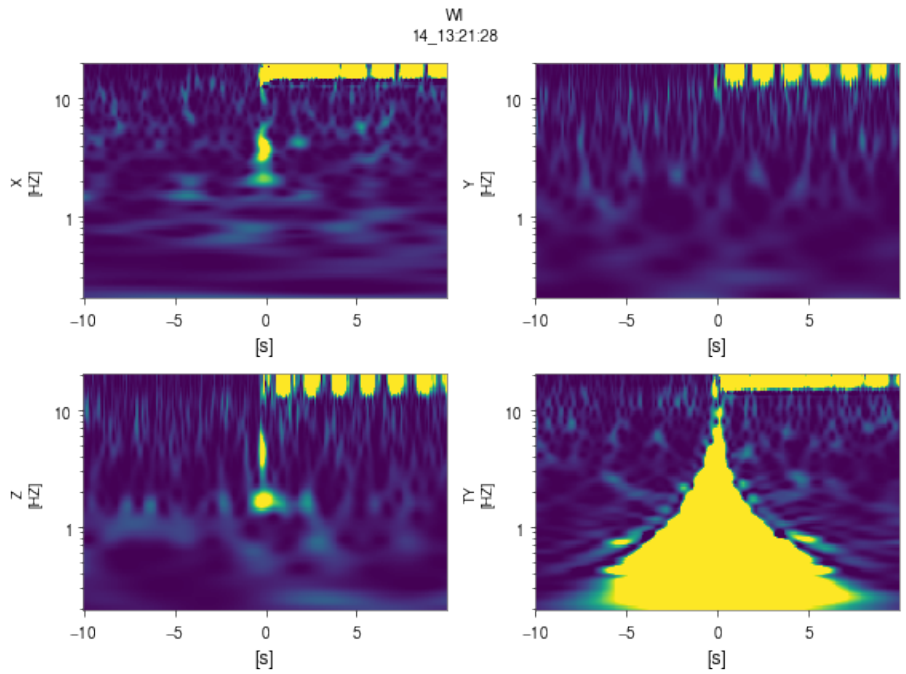

Figure 5.16: Example of one of the witch hat spectrograms. The peculiar feature in θ_y is clear and also the high frequency ring that follows.

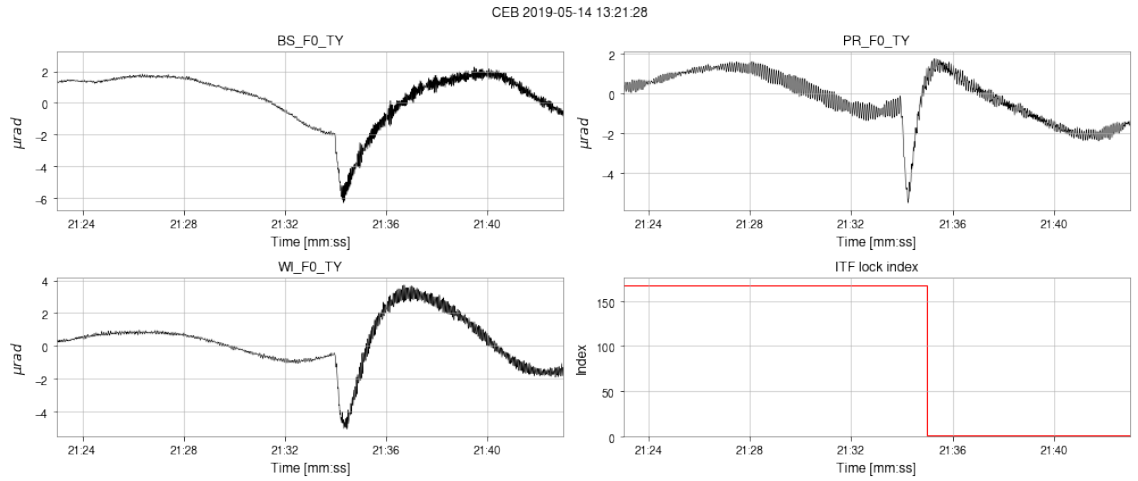

Figure 5.17: Timeseries data for the F0 channel θ_y of 3 different CEB Superattenuators, together with the lock index of the interferometer. All of these events were linked to the witch-hat cluster. The discontinuity seems to be the time-series representation of the "witch hat" feature. This feature appears simultaneously in all of the channels and is in coincidence with the unlock of the instrument.

Tuesday clusters

Some other clusters were found to be linked to the interferometer's bad performance; these were assigned the name "Tuesday clusters". The events that comprise these clusters seem to peak on Tuesday mornings, which coincide with the weekly maintenance, where heavy vehicles and personnel are allowed to operate closer and inside the buildings [45]. This generates high levels of seismic activity and subsequently noise in the suspensions. The cluster shown in fig. 5.18 seems to present loud high frequency noise, like it is presented in the cluster's statistics in fig. 5.19. Usually, these frequencies in seismic excitation get greatly dampened with distance, so normally only local activities create disturbances in this range, which is compatible with the time distribution seen in these clusters. A few experiments were done at the Virgo interferometer by getting close to the Virgo buildings with vehicles and moving near the suspensions, and some links with one of the Tuesday clusters were found. More details will be given in sec. 5.3.

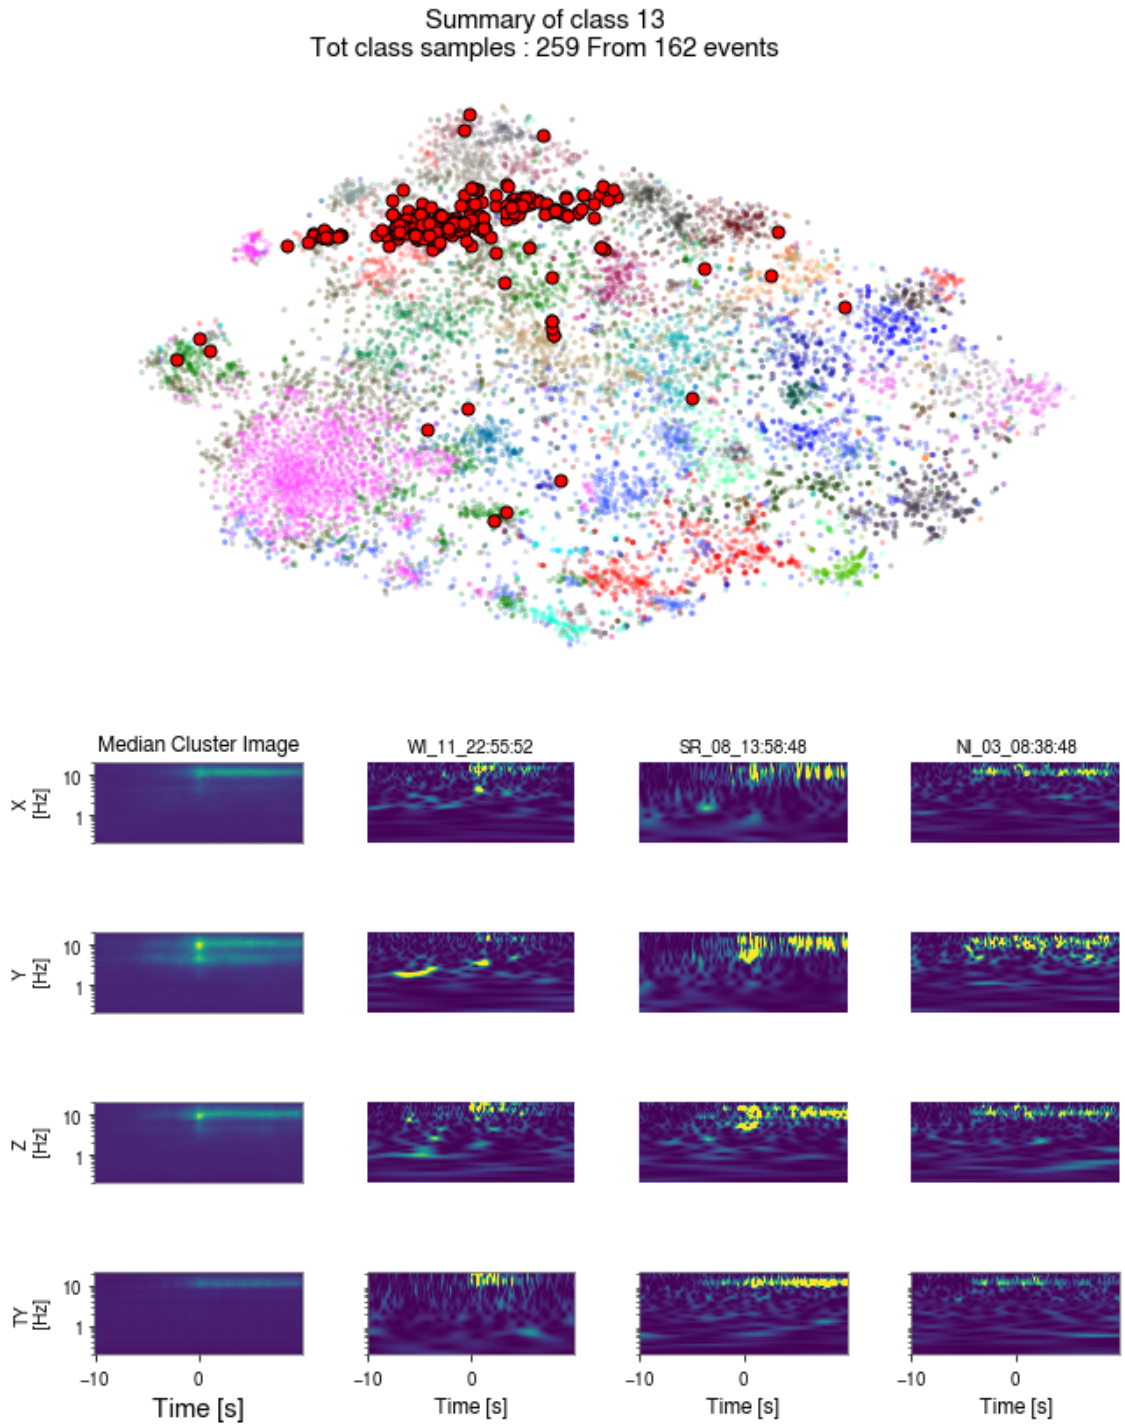

Figure 5.18: t-SNE projection, median spectrogram and examples regarding cluster 13 loud high frequency lines seem to be the most prevalent feature of this cluster.

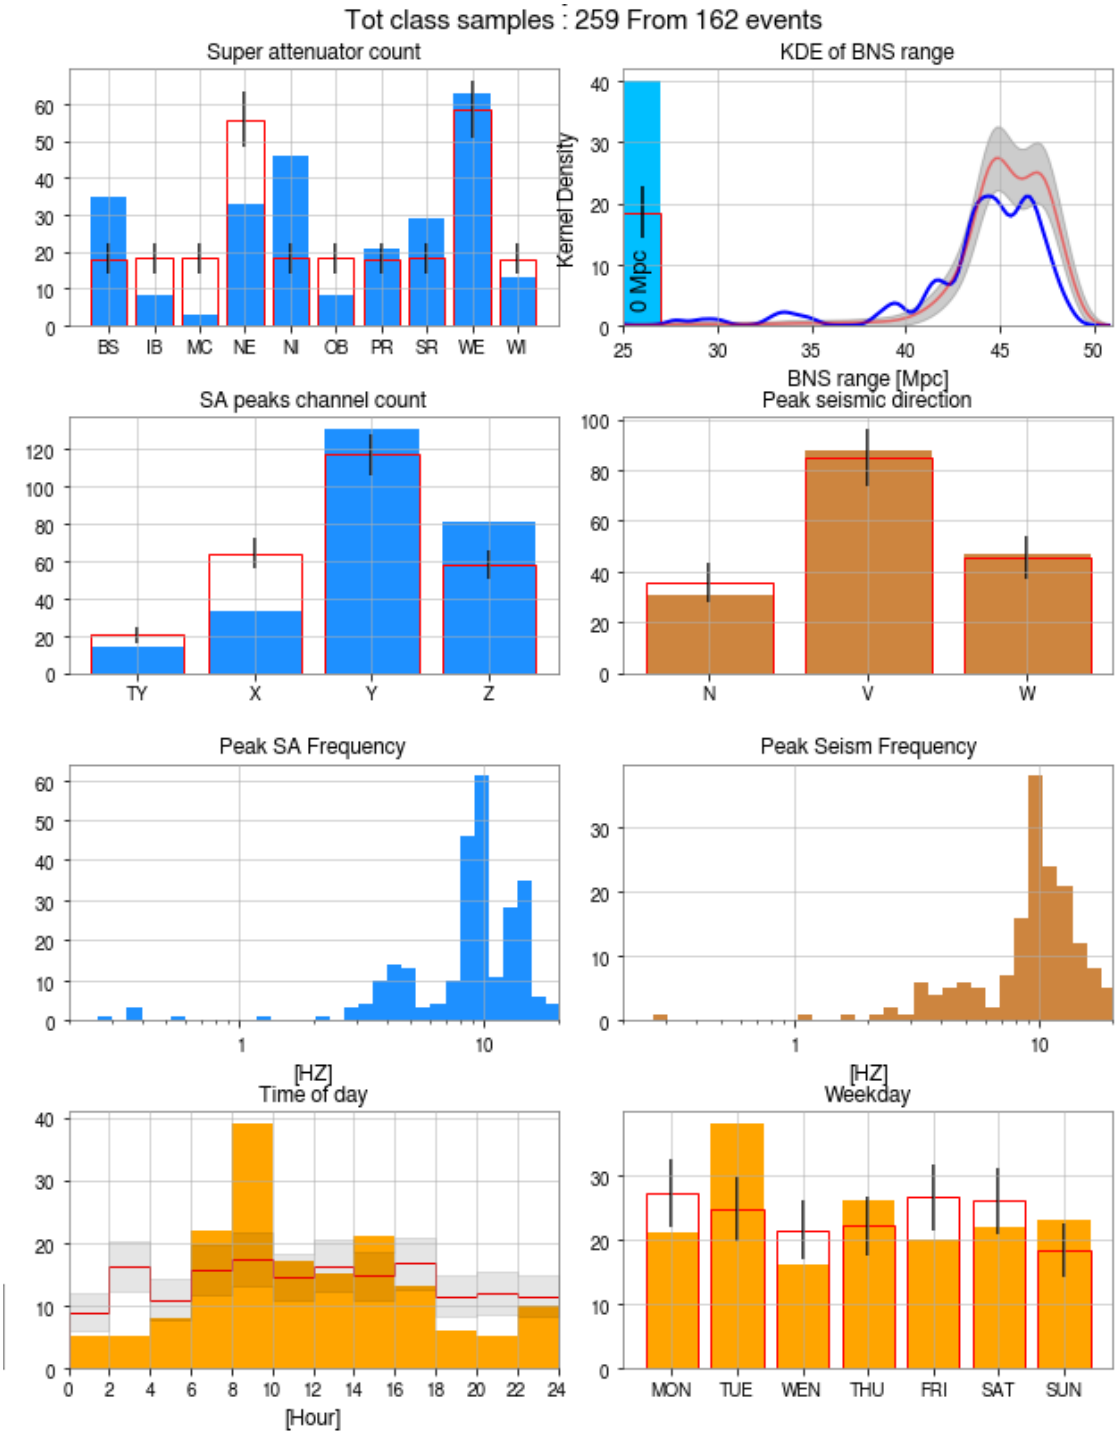

Figure 5.19: Summary statistics of cluster 13. These histograms show the distribution between Superattenuators, the quality of detection, The peak channels (F0 and seismic), the peak frequencies (F0 and seismic), daily and weekly distribution.

Null cluster

Another substantial cluster that is linked to bad data quality is what was called the null cluster, or, the cluster that contains only flat spectrograms. These spectrograms represent white noise only, without any significant features. At a first glance this could seem like a contradiction since one would expect that the cluster with the flattest features would be the one linked to the quietest $h(t)$ data, but what is observed is the exact opposite.

Analyzing the times at which the events that make up the cluster take place, in fig. 5.20, it seems that they peak in coincidence with seismic warnings issued by the SEISMON tool [65] . This algorithm is used to warn the crew that controls the interferometer that a seismic event of considerable strength is happening somewhere in the world and, as the seismic waves resonate around the earth, the seismic activity at the site might degrade the detection quality for extended periods of time [66]. During these periods the interferometer struggles to keep the mirrors in the locked position and the actuators introduce loud non-Gaussian noises in the sensitive band of the detector, sometimes even resulting in an unlock event. Seismic noise induced by distant earthquakes is at extremely low frequencies, usually outside the range of the Q-transforms used in this project. This coincides with the analysis of the BLRMS algorithm frequency range that triggered the acquisition of this data. Most of the events coming from triggers in the range $[0, 0.5]$Hz are actually found in the Null cluster. Even frequencies that manage to enter the correct range still struggle to make a distinctive signal, since to better characterize lower and lower frequencies the Q-transform needs longer and longer time windows. So the lower the frequencies will get the worst representation with respect to the higher ones, given a time window of the same duration. This is probably the main reason why flat spectrograms are linked to low frequency noise and so bad interferometer performance. In future works, it could be useful to extend the study range to better represent the lower frequency domain.

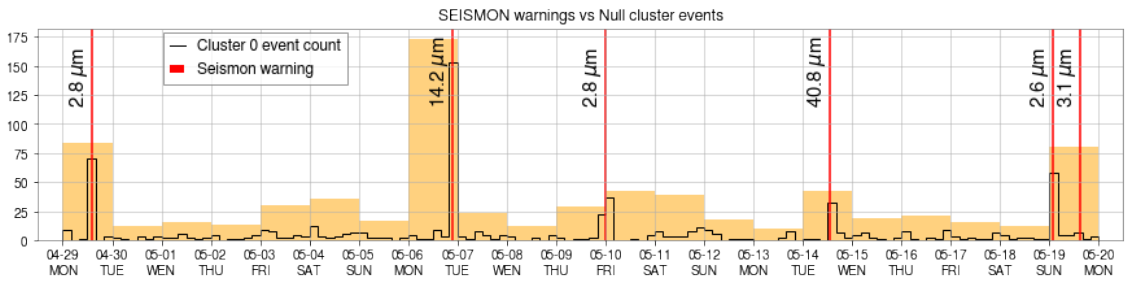

Figure 5.20: Time distribution of the events that make up the null cluster (black line) over imposed with the seismic warnings arrival time (red lines) and expected amplitude of ground motion. The correlation between the two is pretty evident

Although not much useful analysis of this data can be provided due to the difficulty that my algorithm faces in characterizing these low frequencies, it is worth mentioning that by having the algorithm cluster these events together, these will not pollute the analysis of the dataset as a whole. This again shows another possible application of the GMM algorithm, as a dataset cleaning tool: grouping together "not interesting" data it makes the analysis of the "interesting" events much easier.

5.2.4 Correlation with glitches

As we mentioned before, the fast identification of loud transient noises in $h(t)$, the glitches, is of utmost importance for fast gravitational wave searches and improving the detector sensitivity. The aim of this section is to find out if some clusters show some correlations with glitches happening in $h(t)$. The Gravity Spy dataset of Virgo glitches was used as a reference to find labeled glitches in the period of the analysis[32]. As it is possible to see in fig. 4.12 for about a third of the period in question, no glitches were found: the Gravity spy dataset is actually incomplete since it is mostly interested in glitches happening when the interferometer is in full science mode. This could induce some biases in the analysis and for more fine studies more precautions must be taken. Nonetheless, some preliminary investigations can still be made, and some clusters with statistically significant correlations with glitches were found.

The most interesting one seems to be the cluster with number 13, which is one of the clusters peaking on Tuesdays mentioned in the previous section. Some summary statistics of this cluster are reported in fig. 5.18 and 5.19 . The percentage of events that coincide with glitches is just 13% but is more than double the rate predicted by the null hypothesis, which would give just a 6%. By running a Bayesian p-value test, based on the beta function, the 6 % null hypothesis has a p-value of $p = 2.5 \times 10^{-3}$. So the null hypothesis is rejected with 3σ significance. A few examples of glitches that coincide with cluster 13 components are present in fig. 5.21. In this example all of the glitches were labeled as scattered light, which is a class already known for being possibly linked to seismic motion. Another interesting thing is how the glitches appear at frequencies different from the Superattenuaotr excitations, meaning that some kind of non-linear and hard to model behavior might be at play here. This makes cluster 13 a great candidate for future analysis.

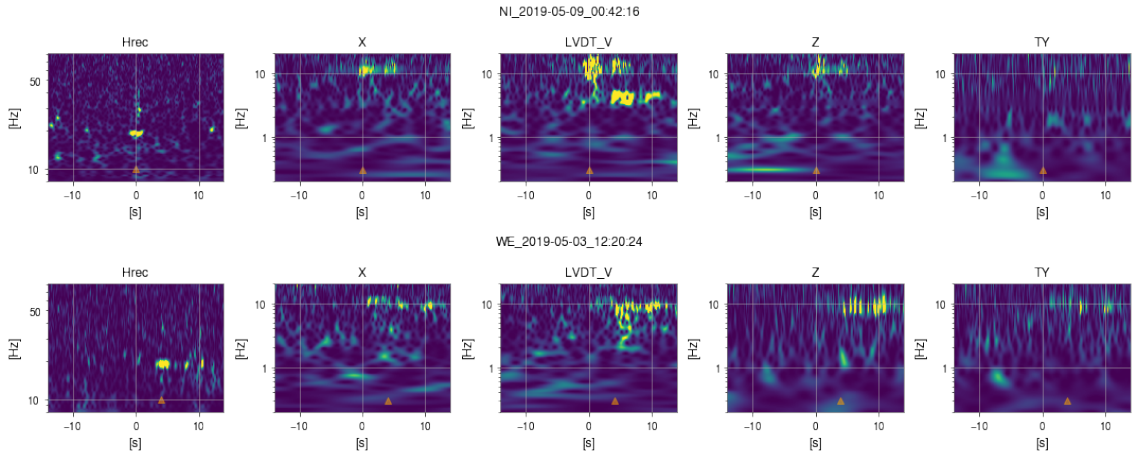

Figure 5.21: 2 examples of spectrograms belonging to cluster 13 where coincidences with glitches were spotted. The time at which the glitch is signaled in the Gravity Spy dataset is noted with an orange arrow to allow better linkage with spectrograms.

5.3 Simulation of anthropic noise at the Virgo site

During the analysis of the performances of the algorithm, environmental channels and comparisons with previous works were used to make assumptions about the origins behind the most peculiar clusters. To test some of these hypotheses, it was decided to make some experiments on the site of the interferometer, with the aim of finding whether human and specifically maintenance activity near the buildings was the actual culprit behind some of the disturbances that the algorithm grouped together.

In particular, the goal was to try and simulate activity on the site by getting close to the terminal buildings with a car, to see if the ground motion induced by the vehicle was able to activate the seismic trigger, and subsequently examine in which cluster the algorithm put these events. The terminal buildings were chosen because of the lower amount of activity happening near them, making the process of identifying the car as a source easier.

Two experiments were made. On the first one, we went with a car in the proximity of both the WEB and the NEB, by approaching each building two times, a few minutes apart, parking the car as close as possible to the Superattenuators. In the second experiment, we went to the WEB two times, about an hour apart, this time entering the building itself on foot to replace a faulty acquisition board. For this last experiment, the activity performed would have been really similar to the one that could happen during maintenance periods of observation runs. A few caveats of these experiments must be addressed:

- These experiments were done almost 4 years after the data on which the algorithm has been trained were taken, so some changes to the Superattenuators in question and the buildings around might have modified how the Superattenuator reacts to the same seismic excitation.

- During both experiments, the detector was in the commissioning phase: during this period the interferometer is unlocked and noise injections are usually run on the detector, which could interfere with the analysis. When the interferometer is locked its output is used in the feedback control loop to keep the mirrors in position. Whenever the detector is down the control is more coarse, thus spectrograms might also have different features with respect to the ones made during runs. Also, the fact that the interferometer was not operating meant that the direct impact of our activity on the detection could not be investigated.

During the first experiment, the seismometers in both buildings detected higher-than-normal seismic activity that coincided with the different arrivals of the car, but it was not strong enough to trigger the download algorithm. This could indicate that light cars moving in the proximity of the buildings might not be the cause behind any of the clusters. It should still be noted that the experiments were done during the day, where there is a higher noise floor in that frequency region due to the large amounts of vehicles moving in the distance, if these experiments were to be run again at night or during holidays the results could be different.

By downloading manually the data, a few interesting features were still found, that are worth mentioning. The first peculiar feature was found by the Superattenuator team which helped me in these experiments. This feature was found in the accelerometer timeseries reported in fig. 5.22 and fig. 5.23. The car's arrival at the Superattenuators leaves a clear signature: there is first a "bump" in the signal and then the measurement settles a bit above 0 for the whole time the car stays in the parking lot. This is most clear by analyzing the data of the x-accelerometer, which is the one that is pointed towards the car. This measurement is compatible with a slight tilt of the tower that gets only partially corrected, since if the F0 platform were inclined, the vertical gravitational acceleration g would creep into the horizontal accelerometers, creating the DC offset from the standard 0 value. Both the frequency needed to represent the bump and the DC offset are outside the range of measurement of the spectrograms, so they could not be studied by my original algorithm. Nonetheless, these are interesting results that point toward the need to develop an algorithm capable of handling the lower frequencies.

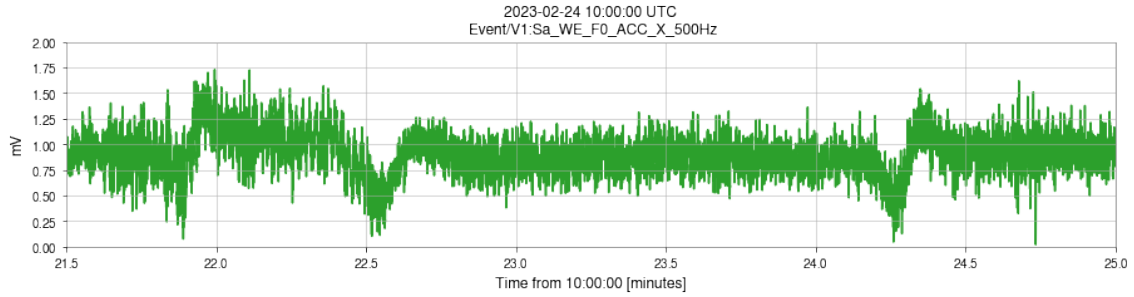

Figure 5.22: accelerometer time-series in the x direction during the car experiment at the WEB. At the times coinciding with the car arrivals and departures, the accelerometer presents peculiar bumps, and after that, the measurement presents a DC offset for the whole time the car stays in the parking lot.

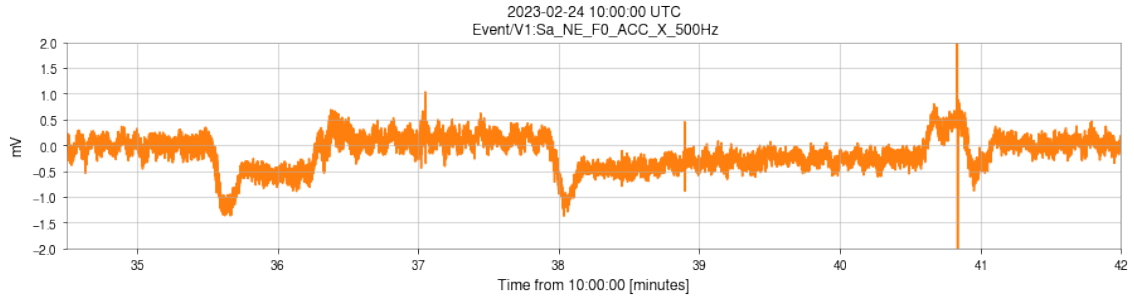

Figure 5.23: accelerometer time-series in the x direction during the car experiment at the NEB. Like in the previous case, the bumps coincide with the car's arrival and departure times, at 35:40 and staying until 36:20 and then coming back again at 38:00, leaving at 41:00. The DC offset stays for the whole time the car is parked.

The second peculiar observation is what is probably the effect of the car engine on the Superattenuator. The spectrograms of both the WEB seismic channels and the Superattenuator x channel around the first and second arrival are shown in fig. 5.24. On the first approach, the car engine was kept running for about a minute before departing again. On the second approach, a few minutes later, the engine was turned off right away. At the time that coincides with the car engine being on and close to the buildings a peculiar line at 24.4 HZ is present. This frequency is consistent with an engine running at $\sim$ 1400 rpm or an overtone of an engine running at $\sim$ 700 rpm, which are typical car engine idling rates. This line lives outside the range of study of the experiment, so it could not be the cause of any of the disturbances found by the main algorithm. Still, it is a hint that in the future it might be interesting to explore in more detail the higher frequency range, where local activity has a clear impact on the F0.

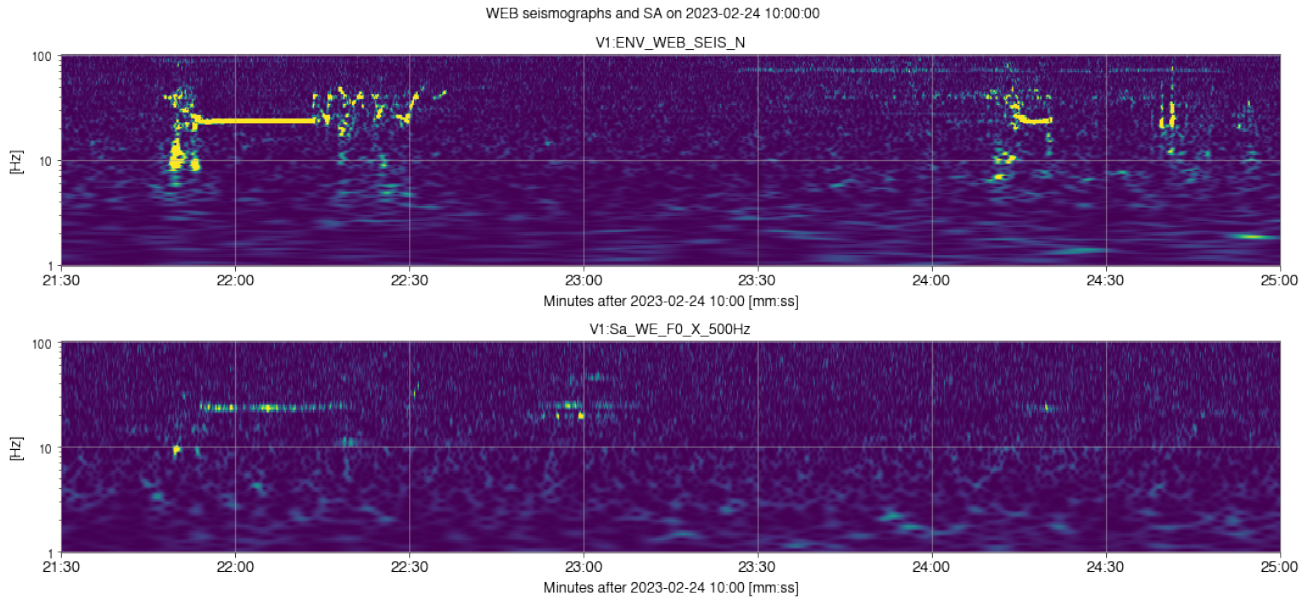

Figure 5.24: Spectrograms showing the Seismometers North and Superattenuator x channels of the WEB around the first and second arrival of the car, at 21:50 and 24:15 minutes after 10 UTC on the 24 of February 2023. The frequency line compatible with the car engine idling at 24 Hz is pretty noticeable in both signals.

The second run of experiments was more successful for the purposes of the VAE algorithm. During the 2 hours over which the experiment took place, the BLRMS algorithm run on the WEB seismometer was triggered 19 times, and quite a few of the events can be confidently traced back to our activities on site. By running the previously trained GMM, these events can be assigned to one of the 40 clusters that have been studied so far. Between the 19 events, the algorithm assigned 4 of them to cluster 13, which is one of the Tuesday clusters that was examined in sec. 5.2.4. These events all coincide with moments when the crew was walking around the building and are characterized by quite a strong signature and high-frequency lines. These are visible in fig. 5.25 that reports these 4 triggers. It was previously speculated that this cluster is linked to local activity happening during maintenance, and this result seems to corroborate these findings. This cluster is also one that presents the strongest correlation with $h(t)$ glitches, as was discussed in sec. 5.2.4 so it would be interesting in the future to analyze in better detail the excitations behind this cluster since they seem to be linked to the quality of detection.

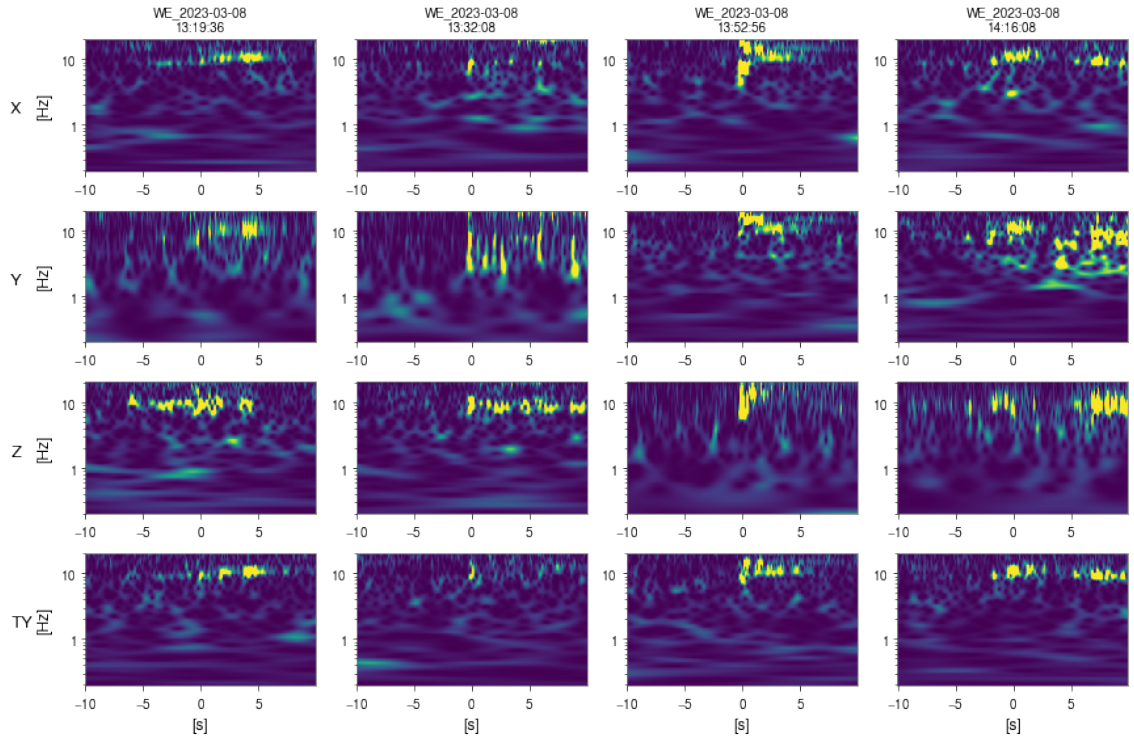

Figure 5.25: Spectrograms that, during the experiment, got put in cluster 13 by the algorithm. All of these seismic excitation occur at times compatible with us walking around the WEB.

5.4 Future applications

A few shortcomings of the algorithm were spotted during the analysis, so it is deemed worthwhile to discuss how future application of this algorithm could improve its performances.

One limitation that really hindered the whole work was the small training dataset, since the 13 000 spectrograms, acquired by analyzing just 3 weeks of data, were deemed not enough for a comprehensive analysis. The main bottleneck in this regard is the time needed for the download of data. It took more than 10 days for the construction of the final dataset, so more resources to download the data might be needed. Another problem linked to the data was the asymmetries in the building seismic count that can be directly linked to how the BLRMS algorithm interacts with the different seismic landscapes. The less sensitivity to low-frequency excitation and the asymmetry between buildings can strongly impact the analysis and insert undesired biases. Modifications of the BLRMS algorithm were tried, like lengthening the running average window and shortening the length of each segment, which showed some promising results.
More channels could be added to the latent space, to better characterize the F0 movements, especially accelerometers. A few tests were done with this data to test the flexibility of the

Ja-Net approach that showed promising results, but were not ready to be presented in this thesis but will be the focus of future work.
Lower and higher frequencies than the range studied both seemed to have interesting features that could be impacting the quality of the detector, so any future application must be equipped to handle these. Due to the disparity of requirements, probably two different algorithms will be deployed, one for the high-frequency range, similar to the one already in place that still deals with spectrograms, and one for the low frequency that deals directly with the time-series.

With the Pisa INFN, we are already looking at how this tool could be implemented during the upcoming O4 observing run, for a quick analysis of the excitation affecting the suspensions of the Interferometer that could help the detector characterization group to identify noise sources and potentially remove them.

Summary

In this chapter, the main results obtained during the thesis were presented. An overview of the performances of the generative model Ja-Net was given, to highlight how this novel VAE architecture's latent space was characterized, by compressing the data more than 2000 times. The rest of the chapter was dedicated to describing the clusters found by the GMM trained on the data compressed by Ja-Net, where the algorithm was found capable of clustering together events originating from the same Superattenuators, proving as a viable way to characterize the instruments and the specific noise that affects them, like the bridges of the FI-PI-LI highway near the NE and the resonant frequencies in the WE. The algorithm was also capable of identifying certain specific clusters linked to a bad detection quality, and in particular, how the "witch hat" cluster is strongly linked to the unlocking of the interferometer. Bad quality was also linked to clusters suspected to be originated from activities on site, like the Tuesday cluster 13, which was found to also be linked to some glitches in $h(t)$ and the fact that it is probably originating from activity on site was validated by experiments done in the vicinity of the buildings. These experiments simulated (and actually performed) usual maintenance activity near the Superattenuators and found the effects that cars and people's movements have on the suspensions. By looking at the raw data from these events it was evidence of how some of the activity performed affected the Superattenuators outside of the frequency range studied by the project, in particular, car engine noise was clearly visible in the high-frequency domain, while the ground tilt due to the weight of the car was way lower. The importance of characterizing low frequencies was also evident by the presence of the Null cluster, which was full of spectrograms devoid of features but still linked to seismic warnings and bad interferometer performance. These failure modes of the algorithm are well known and the last part of the chapter is dedicated to how to improve on them for future implementations.

Conclusions

Gravitational waves detector are complex instruments, that are subject to many kinds of noise that couple together in non-linear and hard-to-predict ways. The state of the interferometer is constantly monitored by an array of more than 100 000 Auxiliary channels. The vast amount of data produced by these streams makes manual analysis impractical but provides a great resource for the training of machine learning models. During this thesis, a deep learning algorithm was developed to study one of the subsystems of the interferometer, the Superattenuator.

After having developed an algorithm capable of detecting higher than normal seismic activity, the channels that describe the Super Attenuator top stage position were downloaded in correspondence to these events. The analysis was run on 3 weeks of O3a data, for a period where the interferometer experienced a diverse set of conditions. After having analyzed the dataset, a custom-made VAE variant was trained on it, to realize a generative model aimed at reducing the dimensionality of the data points by more than 2000 times. After having assessed the performances of the VAE, the GMM algorithm was given the task to find clusters in the now lower-dimensional dataset. The analysis of these patterns yielded some interesting results. Although they all have the same design, it was found that different SuperAttenuators seem to have different responses to noises. The WE was found to present what appears to be a resonance mode at 11 Hz, while the NE seems to be most affected by traffic noises originating from the bridge pillars of a nearby highway. Some clusters were linked directly to unlock events of the interferometer: the effects that the turning off of the feedback control loop has on the suspensions were autonomously recognized by the algorithm. Other clusters were linked to a bad detection quality and their origin is speculated to be on-site activity. These results have been validated by some experiments done in the vicinity of the Virgo buildings, where the algorithm was triggered by the crew movements around the buildings. Some of the excitation produced during the experiment were clustered with events that have been found to have weak, but still statistically significant correlation with glitches in $h(t)$.

This thesis was a first-of-its-kind work to study how machine learning could be used to better understand the behavior of the Advanced Virgo Super Attenuators. This could make us

gain some insights into how these complex systems work and how they can be improved to enhance the detector sensitivity at low frequencies.

The algorithm, although just a proof of concept, showed really promising results. A lot of experience was gained during this thesis work that could improve the capabilities of similar algorithms in the future.

One of the goals for the future of the project will be to improve the algorithm where it showed some weaknesses. One of the main limitations of the current work was the small dataset and so it will be necessary to extend the studied period. Having more data will certainly help build a more complete analysis and will improve the performances of the machine learning algorithms. Rethinking the triggering algorithm could also help in the reduction of the biases present in the current work, by constructing a more balanced dataset between the buildings. Lower and higher frequencies outside the studied range showed possible links with the bad performances of the interferometer, so future algorithms will also need to tackle these ranges. To better characterize the Superattenuator state, accelerometers and different filter channels will need to be included in the analysis. The VAE architecture itself showed some problems due to over-pruning and overfitting, so new architectures will need to be tried

With the team at the Virgo group at the University of Pisa and INFN of Pisa, we are currently looking at ways to implement advanced versions of this algorithm to run during O4, to help the detector characterization team to study the seismic landscape and the performances of the mirror suspension system.

Acknowledgements

Sarà complicato ringraziare tutte le persone che si meriterebbero un posto su questa tesi, ma a me piacciono le sfide.

Ringrazio il prof. Giacomo Ciani per avermi fatto conoscere questa opportunità, e per avermi accompagnato durante tutto il percorso con i suoi consigli.

Ringrazio il prof. Massimiliano Razzano per avermi seguito per tutto questo lungo lavoro e per avermi offerto questa esperienza.

Ringrazio il team del Superattenuatore a Virgo, Valerio Boschi ed Alberto Gennai per aver sopportato le mie domande durante il mio periodo ad EGO.

Volevo ringraziare i membri dell'ufficio 69 per avermi dato asilo e scrivanie nei momenti più duri della scrittura della tesi, grazie Michele, Lucia Maria Antonietta e Diana. Spero tra non molto di potervi chiamare colleghi.

Ringrazio la famiglia Brennero estesa, my brothers Marco, Andrea, Lorenzo, Maria per avermi fatto sentire a casa in in un posto dove non avrei mai creduto di trovare tanto calore. Thanks Ana for teaching me that a person can apparently just survive on cereal and for giving me the skill of changing language in the middle of a frase. Volevo poi ringraziare Gianmarco per essere rimasto mio amico nonostante lo abbia inavvertitamente reso senzatetto per tre mesi.

Grazie a tutta la crew Padovana, che ha reso la mia permanenza in quella città una delle esperienze più belle che abbia mai vissuto. Grazie ai miei ex-coinquilini Rinna e S*m, per l'esagerata quantità di birra che tenevamo nel cucinotto, Rebbe per i pomeriggi passati a imprecare davanti al Bartezzaghi, Mistra per le conversazioni filosofiche sul balcone la sera e Fede per essere stata la migliore bassista della nostra band mai esistita. Grazie ai vicini di casa Rossetti, per avermi sempre dato asilo quando sentivo la mancanza di questa città. Un Grazie particolare a Chiara per aver sempre sopportato la mia totale incapacità di organizzare cose e per la sua pasta alle zucchine. Grazie al resto del gruppo che non può essere nominato, causa ritiro della tesi dagli archivi, Grazie Elena, Marika, Marcello, Michele e Sarbina, per tutti gli spritz e le improbabili serate karaoke. Aspetto con ansia il prossimo Walkshop.

Grazie al Quackmelcorno, per aver reso la triennale l'incredibile esperienza che è stata, e per avermi fatto conoscere lo strano mondo dei giochi di ruolo. Grazie Luchino, Jacopo, GioL,

GioG, Lilli. Non vedo l'ora di rivedere la casa del signor Paolo.

Grazie a tutti i quelli del Fanto condividi musica, per avermi sopportato durante il periodo delle superiori e per sopportarmi tutt'ora quando pacco inesorabilmente le serate giochi da tavolo. Grazie D'Adda, Barro, Carol, Fanto, Verze, Rinna, Pitti, Nico e Batto. Come dicono in Repubblica Ceca, Abajour!

Grazie al gruppo delle bici(s) che conosco da talmente tanto tempo che il termine "amici" mi sembra un po' riduttivo per descrivervi. Grazie Davi, per essere la spalla su cui lamentarsi dell'università, Didi per essere la calma fatta persona io non ho idea di come diavolo fai, Palli, per le vacanze improvvisate che solo tu ti potresti accollare e Riccio, per le serate davanti alla console da DJ. Grazie per essere stati una delle costanti della mia vita, nonstante tutti gli impegni e le distanze.

Grazie alla mia famiglia, che è sempre stata a supportarmi in ogni momento. Grazie a nonno Franco e a nonna Nella, grazie a nonna Marcella e a nonno Vincenzo.

Grazie a mia mamma, Silvia e a mio padre, Amedeo. La tesi è dedicata a voi.
Vi devo tutto.

Bibliography

[1] M. Maggiore, *Gravitational Waves. Volume 1: Theory and Experiments*. Oxford University Press, 2007.

[2] J. Creighton and W. Anderson, *Gravitational-Wave Physics and Astronomy*, ser. Wiley Series in Cosmology. Winheim, Germany: Wiley, 2011.

[3] A. Królak and M. Patil, "The first detection of gravitational waves," https://www.mdpi.com/2218-1997/3/3/59, 2017, last access: 29 April 2023.

[4] S. Weinberg, *Gravitation and cosmology*. Wiley, New york, 1993.

[5] B. P. Abbott *et al.*, "Observation of gravitational waves from a binary black hole merger," *Physical Review Letters*, vol. 116, p. 061102, 2016.

[6] R. Abbott *et al.*, "GWTC-2: compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run," *Physical Review X*, vol. 11, no. 2, p. 021053, 2021.

[7] ——, "GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run," *arXiv e-prints*, 2021.

[8] B. P. Abbott *et al.*, "GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and virgo during the first and second observing runs," *Physical Review X*, vol. 9, no. 3, 2019.

[9] R. Abbott *et al.*, "All-sky search for short gravitational-wave bursts in the third Advanced LIGO and Advanced Virgo run," *Physical Review D*, vol. 104, no. 12, 2021.

[10] D. Radice *et al.*, "Characterizing the gravitational wave signal from core-collapse supernovae," *The Astrophysical Journal*, vol. 876, p. L9, 2019.

[11] C. Cutler, "Gravitational waves from neutron stars with large toroidal b fields," *Physical Review D*, vol. 66, p. 084025, 2002.

[12] K. Riles, "Recent searches for continuous gravitational waves," *Modern ics Letters A*, vol. 32, p. 1730035, 2017.

[13] A. I. Renzini *et al.*, "Stochastic gravitational-wave backgrounds: Current detection efforts and future prospects," *Galaxies*, vol. 10, no. 1, 2022.

[14] M. Armano *et al.*, "The LISA pathfinder mission," *Journal of Physics: Conference Series*, vol. 610, no. 1, p. 012005, 2015.

[15] B. P. Abbott *et al.*, "Advanced LIGO," *Class. and Quantum Grav.*, vol. 32, no. 7, p. 074001, 2015.

[16] F. Acernese *et al.*, "Advanced Virgo: a second-generation interferometric gravitational wave detector," *Class. and Quantum Grav.*, vol. 32, no. 2, p. 024001, 2014.

[17] https://dcc.ligo.org/public/0172/G2002127/018/ObsScen_timeline.pdf, last access: 01 March 2023.

[18] A. Effler *et al.*, "Update on the LIGO, Virgo, and KAGRA detector upgrades for O4," https://dcc.ligo.org/LIGO-G2200736/public, 2022.

[19] https://www.ligo.caltech.edu/MIT/image/ligo20211107a, last access: 01 March 2023.

[20] https://www.mpoweruk.com/figs/Michelson-Interferometer.htm, last access: 10 March 2023.

[21] T. Allocca *et al.*, "Interferometer sensing and control for the Advanced Virgo experiment in the O3 scientific run," *Galaxies*, vol. 8, no. 4, p. 85, 2020.

[22] F. Acernese *et al.*, "Virgo detector characterization and data quality during the O3 run," *arXiv e-prints 2205.01555*, 2022.

[23] ——, "Increasing the astrophysical reach of the Advanced Virgo detector via the application of squeezed vacuum states of light," *Physical Review Letters*, vol. 123, p. 231108, 2019.

[24] S. Braccini *et al.*, "Measurement of the seismic attenuation performance of the VIRGO Superattenuator," *Astroparticle Physics*, vol. 23, no. 6, pp. 557–565, 2005.

[25] L. Trozzo, "Low frequency optimization and performance of Advanced Virgo seismic isolation system," 2018, PhD thesis.

[26] A. Accadia *et al.*, *The Review of scientific instruments*, vol. 82, p. 094502, 2011.

[27] M. Beccaria *et al.*, "Extending the virgo gravitational wave detection band down to a few hz: metal blade springs and magnetic antisprings," *Physics Research Sec. A*, vol. 394, no. 3, pp. 397–408, 1997.

[28] P. Ruggi, "L'attenuazione del rumore sismico nel rivelatore di onde gravitazionali virgo," 2003, PhD thesis.

[29] L. Tianjun *et al.*, "Measurements of mechanical thermal noise and energy dissipation in optical dielectric coatings," *Physical Review D*, vol. 89, 2013.

[30] S. Bahaadini *et al.*, "Machine learning for gravity spy, glitch classification and dataset," *Information Sciences*, vol. 444, p. 172, 2018.

[31] M. Punturo *et al.*, "The einstein telescope: a third-generation gravitational wave observatory," *Class. and Quantum Grav.*, vol. 27, no. 19, p. 194002, 2010.

[32] M. Zevin *et al.*, *Class. and Quantum Grav.*, vol. 34, no. 6, p. 064003, 2017.

[33] M. Razzano *et al.*, https://www.zooniverse.org/projects/reinforce/gwitchhunters, 2022.

[34] T. Accadia *et al.*, "Noise from scattered light in Virgo second science run data," *Class. Quantum Gravity*, vol. 27, no. 19, p. 194011, 2010.

[35] P. Mehta *et al.*, "A high-bias, low-variance introduction to machine learning for physicists," *Physics Reports*, vol. 810, pp. 1–124, 2019.

[36] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.

[37] G. Hinton, "Neural networks for machine learning, lecture 6a," https://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf, 2018.

[38] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv e-prints 1412.6980*, 2017.

[39] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989.

[40] M. Stewart, https://towardsdatascience.com/simple-introduction-to-convolutional-neural-networks-cdf8d307
2019.

[41] E. Cuoco *et al.*, "Enhancing gravitational-wave science with machine learning," *Machine Learning: Science and Technology*, vol. 2, no. 1, p. 011002, 2020.

[42] R. Essik, *Class. and Quantum Grav.*, vol. 30, no. 15, p. 155010, 2013.

[43] https://gcn.gsfc.nasa.gov/gcn3/21509.gcn3, 2017.

[44] N. Mukund *et al.*, "Ground motion prediction at gravitational wave observatories using archival seismic data," *Class. and Quantum Grav.*, vol. 36, no. 8, p. 085005, 2019.

[45] F. Acernese *et al.*, "The virgo o3 run and the impact of the environment," *Class. and Quantum Grav.*, vol. 39, no. 23, p. 235009, 2022.

[46] G. Vajente *et al.*, "Machine-learning nonstationary noise out of gravitational-wave detectors," *Physical Review D*, vol. 101, p. 042003, 2020.

[47] D. P. Kingma and M. Welling, "Auto-Encoding variational bayes," *arXiv e-prints 1312.6114*, 2014.

[48] I. Higgins *et al.*, "beta-VAE: Learning basic visual concepts with a constrained variational framework," *International Conference on Learning Representations*, 2017.

[49] J. Rocca, https://towardsdatascience.com/understanding-variational-autoencoders-vaes-f70510919f73, 2019, last access: 03 February 2023.

[50] S. Samanta *et al.*, "VAE-sim: A novel molecular similarity measure based on a Variational Autoencoder," *Molecules*, vol. 25, p. 3446, 2020.

[51] I. Fiori *et al.*, "The hunt for environmental noise in virgo during the third observing run,"

Galaxies, vol. 8, no. 4, 2020.

[52] S. Kedar and F. H. Webb, "The ocean's seismic hum," *Science*, vol. 307, no. 5710, pp. 682–683, 2005.

[53] G. Saccarotti *et al.*, "Seismic noise by wind farms: a case study from the virgo gravitational wave observatory, Italy," *Bulletin of the Seismological Society of America*, vol. 101, 2011.

[54] I. Fiori *et al.*, https://logbook.virgo-gw.eu/virgo/?r=34368, 2016, entry on the Virgo Logbook.

[55] I. Fiori, L. Holloway, and F. Paoletti, "Studies of the 1-4 hz seism," https://tds.virgo-gw.eu/?content=3&r=1463, pp. 1390–251, 2003.

[56] F. di Renzo *et al.*, "GWdama," https://gwnoisehunt.gitlab.io/gwdama, 2022.

[57] F. Robinet *et al.*, "Omicron: A tool to characterize transient noise in gravitational-wave detectors," *SoftwareX*, vol. 12, p. 100620, 2020.

[58] F. Chollet *et al.*, "Keras," https://github.com/fchollet/keras, 2015.

[59] L. Antelmi *et al.*, "Sparse multi-channel Variational Autoencoder for the joint analysis of heterogeneous data," vol. 97, pp. 302–311, 2019.

[60] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, no. 85, pp. 2825–2830, 2011.

[61] Y. Burda, R. Grosse, and R. Salakhutdinov, "Importance weighted autoencoders," *arXiv e-prints 1509.00519*, 2016.

[62] S. Yeung, A. Kannan, Y. Dauphin, and L. Fei-Fei, "Tackling over-pruning in Variational Autoencoders," *arXiv e-prints 1706.03643*, 2017.

[63] A. Singh and T. Ogunfunmi, "An overview of Variational autoencoders for source separation, finance, and bio-signal applications," *Entropy*, vol. 24, no. 1, 2022.

[64] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," *Jounrnal of Machine Leanring Research*, vol. 9, pp. 2579–2605, 2008.

[65] https://scientists.virgo-gw.eu/Seismon/index_O3.html#O3a, 2020.

[66] M. Coughlinand *et al.*, "Limiting the effects of earthquakes on gravitational-wave interferometers," *CLass. Quantum Grav.*, vol. 34, no. 4, p. 044004, 2017.