
Exploring Long Duration Gravitational-Wave Transients

with Second Generation Detectors

By

MAXIME A. FAYS



School of Physics and Astronomy
CARDIFF UNIVERSITY

A dissertation submitted to Cardiff University in accordance with the requirements of the degree of DOCTOR OF PHILOSOPHY in the School of Physics and Astronomy

SEPTEMBER 2017

ABSTRACT

Minute-long gravitational-wave (GW) transients are currently a little-explored regime, mainly due to a lack of robust models. As searches for long-duration GW transients must rely on minimal assumptions about the signal properties, they are also sensitive to GWs emitted from unpredicted sources. The detection of such sources offers exciting and strong potential for new science.

Because of the large parameter space covered, all-sky long-duration transient searches require model-independent processing and fast analysis techniques. For my PhD thesis, I integrated a set of fast cross-correlation routines in the spherical harmonic domain (SphRad) [50] into X-pipeline [95], a targeted GW search pipeline commonly used to search for GW counterparts of short and long duration GRBs & core-collapse supernovae.

Spherical harmonic decomposition allows for the sky position dependency of the coherent analysis to be isolated from the data [40] and cached for re-use, saving both time and processing units. Moreover, the spherical harmonic approach offers a fundamentally different view of the data, allowing for new possibilities for rejecting non-Gaussian background noise that could be mistaken for a GW signal.

The combined search pipeline, X-SphRad, underwent a thorough internal review within the LIGO collaboration, which I led. The pipeline good functioning was assessed by rigorous tests including comparing a test data set with a standard sky grid-based analysis.

I have developed a novel pixel clustering method that does not depend on the amplitude of potential signals. By using an edge detection algorithm, I quantify each pixel in the spectrogram by its similarity with its neighbours then extract features of sharply changing intensity (or ‘edge’). The method has shown promising results in preliminary tests. A simplified version of the algorithm was implemented in X-SphRad and large-scale testings are currently being processed.

DEDICATION AND ACKNOWLEDGEMENTS

What good timing to do a PhD in gravitational physics! The last few years have been fantastic, with the first direct detection of a gravitational wave... and more since then. I feel incredibly lucky to have taken part in such a unique experience, and proud to have made a modest contribution to the field.

This marks the end of a 8 year long journey. Starting from Belgium with a business degree, I wanted to see the world, learn English and work towards a goal I was passionate about. Well, I did... and what a ride! I lived in Varkaus, St Helens, Aberystwyth, Boise, Svalbard, Yorba Linda, Cardiff and Baton Rouge... and I absolutely love what I do. Home run!

First and foremost, I'd like to thanks my parents for their unconditional support. When first I said I wanted to start a new degree in Physics without any physics background, they thought the worst case scenario was that I'd spend a year in the UK and learn to speak English. It went rather well! A million thanks to Fred & Syl. Un tout grand merci aux "tuteurs" ! Now, we really have to talk about this 'math teacher' job.

A big thanks to Valeriu, I couldn't have done it without you buddy! Good Science is done around good friends with good beers (brewskies!), and we certainly did our fair share of 'Sciencing' together. Cheers for being there during the good times and the bad, you're one of a kind.

A huge, huge thank you to Ruth. Thank you so much for the incommensurable help during the whole writing process. You prepared so many delicious foods, cleaned the house so often and spent so much time fixing my Latex issues. I'm sure I'd still be yelling at my laptop if it wasn't for you! I bet you are as happy as I am that this is over, you won't have to hear me repeating 'and I still need to do this... and I need to finish this, and there is also this...' more times than any sane person would. Next year is my turn to pamper you!

Obviously, I want to thank Patrick. Thanks for all the help you provided me, from the beginning until the very end with a last-minute meeting mere hours before the submission deadline of this document. Thanks for looking out for me during the tough times, and for stepping up when needed. Cheers Boss!

ABSTRACT

I also want to thank people that were instrumental in me taking and following this path. Jean-Luc Dighaye, that invited me in his house 10 years ago to share his passion and opened the world of astronomy. Martin Wilding, for giving a chance to a young Belgian guy with no prior physics background. Daryl Macomb, who gave more freedom than many graduates and was an example showing that both teaching and research was fun... and addictive. Balazs Pinter, for writing the nicest reference letter I have ever read and helping me get in the Cardiff PhD program. Finally, Mario Bisi that spent way more time talking with an undergraduate about PhD application than a post-doc should, and for introducing me to the world of Gravitational Waves.

Finally, I'm sending all my good luck (I've had enough!) to my little brother Boris finishing his PhD. You chose the dark side, finance, instead of advancing Science... so I'm counting on you to make enough money to take care of my expensive lifestyle. Good luck frerot.

Elle est pas Belge la vie ?

AUTHOR'S DECLARATION

I declare this work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

SIGNED: DATE: 26.03.2018



Statement 1

This thesis is being submitted in partial fulfilment of the requirements for the degree of PhD.

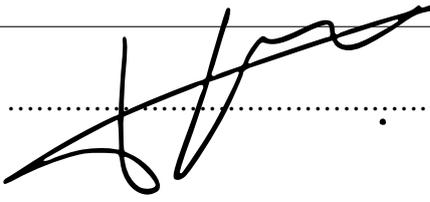
SIGNED: DATE: 26.03.2018



Statement 2

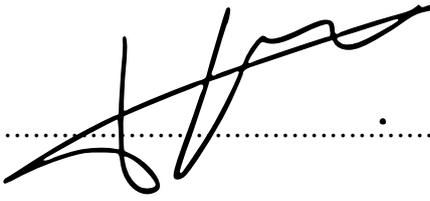
This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references. The views expressed are my own.

ABSTRACT

SIGNED:  DATE: 26.03.2018

Statement 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

SIGNED:  DATE: 26.03.2018

Statement 4

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loans after expiry of a bar on access previously approved by the Academic Standards & Quality Committee.

SIGNED:  DATE: 26.03.2018

TABLE OF CONTENTS

Abstract	i
	Page
List of Tables	xi
List of Figures	xiii
1 Introduction to General Relativity	1
1.1 Einstein's Equations	3
1.1.1 The Metric	3
1.1.2 Ricci Curvature Tensor	5
1.1.3 Stress-Energy	6
1.2 Gravitational Waves	6
1.2.1 Properties of Gravitational Waves	8
1.3 LIGO Detectors	9
1.3.1 Noise Sources	12
1.4 Astrophysical Sources of Long-Duration GW Transients	14
2 Gravitational wave detection by spherical radiometry	17
2.1 Introduction	17

TABLE OF CONTENTS

2.2	Gravitational Wave Analysis	19
2.3	Formalism	20
2.4	Pipeline Workflow	25
2.4.1	Data Pre-Conditioning	25
2.4.2	Sky Position Dependency	28
2.4.3	Time-Frequency Maps	29
2.5	Clustering	31
2.5.1	Background Estimation	32
2.5.2	Software Injections	33
2.5.3	Coherent Cuts and Tuning	34
3	O1 Long-Duration Burst Search	37
3.1	Overview	37
3.2	Observing Run 1	38
3.2.1	Analysed Science Segments	39
3.2.2	Sensitivity	40
3.2.3	Glitches	43
3.2.4	Data Quality	48
3.2.5	Additional Vetoes	49
3.3	Search Parameters	49
3.4	Clustering	50
3.5	Background Estimation Procedure	54
3.5.1	Properties of Background Triggers	58
3.6	On-Source Triggers	64
3.7	Search Sensitivity	69

3.7.1	Astrophysical Models	70
3.7.2	Ad-hoc Models	74
3.8	Upper Limits and Pipeline Comparison	84
4	Enhancement to long-duration burst searches	89
4.1	Introduction	89
4.2	Clustering the time-frequency map	91
4.3	Edge detection	94
4.4	Difference of Gaussians	97
4.5	Preliminary results	98
5	Conclusion	107
6	Appendix A	109
	Bibliography	127

LIST OF TABLES

TABLE	Page
3.1 Breakdown of dataset total duration and science mode time after applying data quality requirements and coincidence. The time processed by X-SphRad is slightly less than the H1-L1 coincident time due to the pipeline requiring a minimum of 512 s of continuous data for analysis.	41
3.2 Channel names and corresponding deadtime for additional vetoes applied specifically in the long-duration O1 searches. The additional vetoes were suggested by the STAMP-AS team [46]. The burst CAT2 short duration vetoes are also effective against RF45 glitches, but have a much larger deadtime ($\sim 2\%$) that was negatively affecting the O1 long-duration Zebragard search. Source: private email conversation.	50
3.3 Top 5 loudest surviving triggers in background.	64
3.4 Surviving foreground candidate events. All were consistent with the background.	68

3.5	Time-frequency properties for astrophysical waveforms. The ADI family covers the lower end of the targeted frequency space with ~minute long signals, while the magnetar family covers the middle and higher end with longer duration. Individual TFmaps of each waveform can be found in Appendix B.	71
3.6	Detectability of the astrophysical waveforms. The optimal SNR is the SNR needed to recover 50% of the waveforms in a single aLIGO detector in noise estimated from aLIGO design sensitivity, weighted by a factor of 3 to match current aLIGO sensitivity. For this estimation, the waveforms were generated assuming face-on orientation. Taken from [29]	73
3.7	Properties of interest for simple ad-hoc waveforms.	77
3.8	90% confidence upper-limits for the astrophysical waveforms.	88

LIST OF FIGURES

FIGURE	Page
1.1 Illustration of the effect of a GW passing along the z direction through a ring of particles on the xy plane. The wave's plus polarisation h_+ and cross polarisation h_\times cause an alternating squeezing and stretching of the test masses, offset by 45 degrees. Illustration taken from [27]	10
1.2 Simplified illustration of one of the LIGO detectors. A laser beam is split into two components, then travels to mirrors in perpendicular arms. In cases where the two arms are exactly the same length, the re-combined laser will have destructive interference at the photodetector, and therefore cancel out.	11
1.3 Sensitivity of the iLIGO detectors. The quantum noise is the limiting factor at frequencies > 100 Hz. The detector is limited by the other degrees of freedom below 100 Hz, which includes instrumental cross couplings. The strong line features are caused by instruments and calibration lines. Taken from [13]	12

1.4 : Spectral variance of the advanced LIGO detectors during ~1 hour around GW150914 for the Livingston detector (L1). The sensitivity of the detector typically fluctuates by a factor of 5, with the more noisy frequency bands exhibiting variance of few order of magnitude (see the 500 Hz band for example.) 15

2.1 Workflow of the X-SphRad pipeline. X-pipeline is used for the pre- and post- processing, while the spherical radiometer SphRad is use as core detection engine. 26

2.2 Skymap of a network of 3 detectors. Each pair of detectors produce a ring of constant delay where the crossing of the rings is the most probable sky location of the signal.Taken from [50] 30

2.3 Comparison of time-frequency map using the full (Eq. 2.17) (left) and reduced (Eq. 2.18) statistic (right). The two agree well, as no difference is visible between features of both maps. Taken from [50] 31

2.4 Sky map energy E_{sh} versus incoherent energy I_{sh} for simulated data (black) and injections of various amplitudes (coloured). The coloured crosses on top of the background are due to the post-processing misselecting background triggers as injected signals. Taken from [50] 35

-
- 3.1 Representation of aLIGO observing time during O1 for both the H1 and L1 detectors, as well as coincident observing time from September 12, 2015 to January 19, 2016. H1 observed in science mode during 78.8 days (62% duty cycle) and L1 during 67.2 days (55% duty cycle). The total network duty cycle for science mode is 40%. The pictured segments do not take into account subsequent data quality vetoes. See also Table 3.1. Generated with [76] 40
- 3.2 Representative noise amplitude spectrum during the O1 period for H1 (red), with aLIGO ideal sensitivity for O1 configuration and 25W laser power superimposed (black). The solid traces represent the median sensitivity and the shaded regions indicate the 5th and 95th percentile over the analysis period. The narrowband features in the spectra are due to known mechanical resonances, mains power harmonics, and injected signals used for calibration [12]. Non-stationarity in the detector noise manifests as variations in the level and shape of these sensitivity curves over time. The grey curve represents the design noise budget for aLIGO during O1. Snapshot taken from November 12, 2015. Generated with [76] 41

3.3 Representative noise amplitude spectrum during the O1 period for L1 (green), with aLIGO ideal sensitivity for O1 configuration and 25W laser power superimposed (black). The solid traces represent the median sensitivity and the shaded regions indicate the 5th and 95th percentile over the analysis period. The narrowband features in the spectra are due to known mechanical resonances, mains power harmonics, and injected signals used for calibration [12]. Non-stationarity in the detector noise manifests as variations in the level and shape of these sensitivity curves over time. The grey curve represents the design noise budget for aLIGO during O1. Snapshot taken from November 12, 2015. Generated with [76] 42

3.4 Average detection range for binary neutron star mergers using template-based filtering during O1, at H1 and L1. H1 is consistently more sensitive than L1, with a detectable range around 10 Mpc higher. 43

3.5 Time-frequency map showing the wandering line glitch at ~ 620 Hz. The associated pixels are the loudest connected pixels in the block time. Noisy pixels of comparable amplitude around the wandering line have a tendency to get clustered together and corrupt neighbouring frequencies. A frequency notch between 590 Hz and 640 Hz was applied to remove the effect of this glitch family. 45

-
- 3.6 Time-frequency map illustrating the effect of glitches on cluster selection. The time-frequency map contains 3 particularly loud glitches at 200 s, 260 s and 360 s as well as a simulated signal from 150 s to 250 s in the 700-900 Hz band. The red lines indicate the bounding boxes of the triggers followed up by X-SphRad. All 30 of the selected triggers are part of the glitches, despite the simulated signal being loud enough to be visible by eye in the data. The white band at ~ 500 Hz is caused by the imperfect whitening of the instrumental violin modes. White grid lines were added to the figure. 46
- 3.7 Time-frequency map showing the effect of glitches on candidate event reconstruction. Two glitches at 20 s and 260 s, as well as multiple brief broadband glitches at 50 s, 100 s and 400 s (below 100 Hz) are present in the data. There is also a downward-chirping simulated signal from 350 s to 410 s. The red lines indicate the bounding boxes of the triggers followed up by X-SphRad. The short glitch at 400 s gets clustered with the simulated signal, resulting in an overestimation of the signal bandwidth and corruption of the estimated energies. 47

3.8 Time used to generate the delay matrix and to process 64 seconds of data, as a function of spherical harmonic number l_{max} . X-SphRad has a fast processing time but a significantly longer startup time. The dominant cost is the background analysis using internal timeslides. Each 512 s background job processes 165 circular timeslides (see Section 3.5 for details). With the chosen parameters $l_{max} = 60$, the estimated startup time is approximately 1 minute, compared to about 200 minutes for processing the internal timeslides. The startup time does not significantly affect the total computational cost. Taken from [50] 51

3.9 Example of long duration signal being fractured into multiple sub-triggers. The signal is split into 8 distinct clusters, each containing only a fraction of the signal total SNR, and will therefore have sub-optimal energy and parameter estimations. The simulated signal is also segmented around 120 Hz due to the effect of the whitening filter. 53

3.10 Effect of the smart clustering on the 50 % distance efficiency recovery for the astrophysical ADI waveforms. The smart clustering increases the recovery distance by a factor of 2 for Adi-D to 3 for Adi-B. 54

3.11 Histogram of trigger rate versus significance threshold, including triggers from all upper-limit background jobs and triggers from all on-source jobs, before and after all vetoes & background rejection cuts. There is a clear correlation between the off-source and on-source distributions before cuts, as discussed in the text. 57

3.12	Skellam test for the off-source triggers surviving all cuts comparing the empirical distribution of $N_{i+1} - N_i$ to the Skellam prediction. The empirical distribution of the background triggers after cuts in blue is consistent with the Skellam distribution in yellow.	58
3.13	Background and foreground candidate events E_{sh}/I_{sh} energy plane. Low energy candidate events (< 10) tend to have larger sky map energy E_{sh} than incoherent I_{sh} energy. This observation motivates the alpha cut (Eq. 2.24).	59
3.14	Distribution of background triggers before cuts, after the linear and alpha cut as well as final distribution after all cuts vs peak frequency. Most triggers are below 300 Hz before cuts. The linear cut does not seem to discriminate based on the frequency, as the distribution post linear cut follows the original distribution. The alpha cut exhibits a frequency-biased selection effect with most of the surviving triggers located above 600 Hz. The final distribution after all cuts follows the alpha cut distribution, indicating the cut to be the most restrictive overall.	60
3.15	Alpha ratio (Eq. 2.24) vs peak frequency for background candidates. The vast majority of low frequency triggers tend to have a low alpha ratio and are thus more likely to be rejected by the alpha cut.	61
3.16	Correlation matrix between properties of background triggers before all cuts. The number of pixels is correlated with the duration and bandwidth, themselves correlated with the trigger peak frequency. The effect of the weighting factor during the clustering phase are seen with the relatively high correlation between number of pixels, bandwidth and duration with the sky map energy E_{sh} and incoherent energy I_{sh}	62

3.17 Correlation matrix between properties of background triggers after all cuts. The number of pixels is correlated with the duration and bandwidth, themselves correlated with the trigger peak frequency. The effect of the weighting factor during the clustering phase are seen with the relatively high correlation between number of pixels, bandwidth and duration with the circ energy and circ inc. 63

3.18 Loudest off-source trigger 65

3.19 Zoom 66

3.20 Surviving background (black) and foreground (red) triggers during the O1 period, after coherent cuts are applied. The triggers are ranked by significance. Foreground triggers are spread throughout the O1 period, with a slight concentration over the last few months of the run. All foreground triggers are consistent with the background. 67

3.21 Loudest surviving on-source event 68

3.22 Illustration of the parameter space covered by the full set of simulated signals used during the O1 analysis. Each shaded rectangle represents the duration and frequency range of one of the 16 waveform types simulated. The whole frequency space is covered with a slight concentration at lower frequencies in aLIGO’s most sensitive band. 70

3.23 Comparison of 50% efficiency and 90% sensitivity distances of X-SphRad during O1 for astrophysical waveforms. Differences between ADI waveforms can be explained by the spread of power in longer durations and wider bandwidths. The differences between the ADI family and the magnetar family are due to the magnetar lower intrinsic amplitudes. . . 72

-
- 3.24 Energy spread of ADI-E in time and frequency. The waveform energy over time is slowly decreasing, while the spread of energy in frequency shows an inverse chirp with power concentrated towards the end of the waveform. This is due to the monochromatic-like behaviour of the waveform at the end of its lifetime. 74
- 3.25 Recovered bandwidth for ADI-E: evolution of mean recovered bandwidth per distance from source. The shaded region corresponds to 1-sigma standard deviation. The dashed line shows the exact bandwidth of the waveform model. The waveform bandwidth is likely to be recovered accurately starting from 9 Mpc. The effects of Fourier leakage are observed up to 5 Mpc and artificially inflate the bandwidth recovery by saturating adjacent pixels, subsequently they are all clustered together. 75
- 3.26 Recovered duration for ADI-E: evolution of mean recovered duration per distance from source. The shaded region corresponds to 1-sigma standard deviation. The dashed line shows the exact duration of the waveform model. The waveform duration is accurately measured from 5 Mpc onward. 76
- 3.27 Mean number of pixels per trigger for each injection scale for ADI-E. The shaded region corresponds to a 1-sigma standard deviation. 76
- 3.28 Energy spread of Magnetar-F in time and frequency. The waveform energy is decreasing over time with maximum at the beginning of the waveform. The spread of energy in the frequency domain is relatively constant, with a faster decrease towards the end of the waveform. . . . 77

3.29 Recovered bandwidth for Magnetar-F: evolution of mean recovered bandwidth per distance from source. The shaded region corresponds to a 1-sigma standard deviation. The dashed line shows the exact bandwidth of the waveform model. The waveform bandwidth is never fully recovered and falls exponentially after 1 Mpc. 78

3.30 Recovered duration for Magnetar-F: evolution of mean recovered duration per distance from source. The shaded region corresponds to 1-sigma standard deviation. The dashed line shows the exact duration of the waveform model. The waveform duration is never fully recovered and fall exponentially after 1 Mpc 79

3.31 Effective distance for the simple ad-hoc waveforms during the O1 period at a FAR of 1 in ~ 50 years. 79

3.32 Low amplitude software injection of MONO-A waveform. The waveform is visible at ~ 90 Hz, from 0 to 150s. The effect of the zero-phase linear predictor error filter is visible from 150 s onward, where a clear line of less-than-average power is present at the same frequency. The effect of the filter has been enhanced using a gaussian average window over 51 seconds. 80

3.33 High amplitude software injection of MONO-C waveform. The waveform can be seen at 405 Hz, between 150 to 400 s. The effect of the zero-phase linear predictor error filters are seen up to 150 s and after 400 s onward, where a clear line of less-than-average power is present at the same frequency. The effect of the filter has been enhanced using a gaussian average window over 11 seconds 81

-
- 3.34 High amplitude software injection of Line-A waveform. Potential GW candidates selected by the pipeline are denoted by a black rectangle corresponding to the event bounding box. The waveform can be seen at $\sim 50\text{Hz}$ to 200 Hz , from 60 to 300 s . Despite the high amplitude of the injection, the pipeline split the waveform at 80 s , and clusters glitches (at 90 and 200s). The wandering line glitch at 630 Hz is also selected as potential GW candidate, as well as a wide-band glitch at 340s 82
- 3.35 Energy spread of Line-B in time and frequency. The waveform energy over time is relatively constant, with a hann window taper at the beginning and end. The spread of energy in the frequency domain shows an inverse chirp with power concentrated in the end-life of the waveform. 83
- 3.36 High amplitude software injection of the Quad-A waveform. The waveform can be seen starting from 330 s to 360 s , with the bounding boxes selected by the pipeline in red. A loud line glitch is visible at $\sim 10\text{ Hz}$. . . 84
- 3.37 High amplitude software injection of the Quad-B waveform. The waveform can be seen starting from 30 s to 90 s , with the bounding boxes selected by the pipeline in red. A loud line glitch is visible at $\sim 10\text{ Hz}$. . . 85
- 3.38 Comparison of the whitening filter effect on a sine-Gaussian waveform. The overly loud signal (top) will corrupt the filter and downweight the values of the other pixels in the frequency bins. For moderate amplitude, the start and end of the waveform are discarded due to the Gaussian envelope (mid). At lower amplitude (bottom) the whitening filter effectively removes the sine-Gaussian. 86

-
- 4.5 Effect of the DoG on a time-frequency map generated from data from the LIGO O1 period, with different kernel sizes and an added simulated signal. The differences due to kernel sizes are apparent in the excess of noise in the bottom and top 40 Hz of the time-frequency maps, and on the simulated signal. top-left) $m=2, n=7$ top-right) $m=2, n=12$ bottom-left) $m=5, n=7$ bottom-right) $m=2, n=3$. All time-frequency maps were thresholded by removing pixels with value lower than zero. For production mode, we recommend parameters with a ratio following Eq. 4.4 101
- 4.6 Comparison of the current X-SphRad clustering with the DoG for the ADI-E waveform. The DoG method was consistently more sensitive than the current clustering implementation in X-SphRad at distance above 500 Mpc. The DoG recovers about twice as many pixels per candidate event as X-SphRad at the same distance. More information as discussed in text. 104
- 4.7 Comparison of the current X-SphRad clustering with the DoG for the ADI-D waveform. The DoG method was consistently more sensitive than the current clustering implementation in X-SphRad at distance above 100 Mpc. The DoG recovers about twice as many pixels per candidate event as X-SphRad at the same distance. More information as discussed in text. 105
- 6.1 ADI-A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 110

6.2 ADI-B: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 111

6.3 ADI-C: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 112

6.4 ADI-D: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 113

6.5 ADI-E: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 114

- 6.6 Magnetar F: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 115
- 6.7 Magnetar G: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 116
- 6.8 Quad-A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 117
- 6.9 Quad-B: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 118

6.10 SG A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 119

6.11 SG C: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 120

6.12 WNB A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 121

6.13 WNB B: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 122

- 6.14 LINE A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 123
- 6.15 Line B: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 124
- 6.16 Mono A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 125
- 6.17 Mono C: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit 126

INTRODUCTION TO GENERAL RELATIVITY

Gravitational Waves (GWs) are direct consequences of Einstein's General Theory of Relativity (GR) [52], introduced in 1916. The existence of GWs was indirectly confirmed 60 years later by Hulse & Taylor with the binary pulsar PSR1913+16 [65].

Direct detection of GWs is a challenging task. As their amplitudes when passing through Earth-based detectors are smaller than the size of a proton [35], major scientific innovations had to take place before having the technical capability of picking up GWs. The initial Laser Interferometer for Gravitational wave Observatory (iLIGO) detectors described in [44] underwent regular upgrades over the last two decades with ever increasing sensitivity. However, no GW was detected during the first six iLIGO science runs which included joint observation time with the French-Italian Virgo detector. Large-scale improvements to the LIGO detector lowered the noise floor in the target frequency band by a factor of 10, at which stage the detectors were renamed Advanced LIGO (aLIGO) [12].

In 2015, nearly a hundred years after their original prediction, aLIGO recorded the first GW signal [13]. Originating from a binary black hole (BBH) system, the signal's detection made history by opening a new window in astronomy with the first:

1. detection of a black hole (BH),
2. detection of a binary black hole system,
3. detection of a binary black hole merger.

Since then, LIGO recorded 3 such events indicating that there is a population of heavy black holes that can form in nature, with binary black-hole systems forming and merging within the age of the universe at a detectable rate [13, 14].

Cosmological models not predicting heavy BHs or low BBH merger rates are already starting to be constrained. Future observations of different systems will bring incredible insights to the universe, the rules of nature in the strong gravitational field regime, and the inner physics of nature's most energetic events that cannot be observed in the electromagnetic spectrum.

The next few years will mark the emergence of GW astronomy as a field.

This document aims at giving an overview of the current status of the gravitational physics field, with a particular focus on searches targeting GW signals lasting between 10 and 1000s, with minimal assumptions about the characteristics of the signals.

After having introduced necessary materials about the generation of GWs and the LIGO detectors, we will describe the X-SPhRad All-sky Long-duration GW transient

search carried out on the first aLIGO run (O1). Moreover, we will suggest alternative approaches to improve the search sensitivity.

This chapter aims to give a brief introduction to Einstein's General Relativity theory and give working key concepts used in this work. It is adapted from [72, 86, 56], and more rigorous approaches can be found in most textbooks, for example [61, 77].

1.1 Einstein's Equations

General relativity describes spacetime as a four-dimensional manifold, globally non-Euclidean and locally Minkowskian. This property allows the manifold to be completely described using the familiar Euclidean spaces.

Spacetime geometry is defined by the Einstein tensor $G_{\mu\nu}$ and is determined by the distribution of mass-energy $T_{\mu\nu}$, the stress-energy tensor, representing the amounts of energy, momentum, pressure, and stress in the space. We see that

$$(1.1) \quad G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{8\pi G}{c^4}T_{\mu\nu},$$

where $g_{\mu\nu}$ is the metric, $R_{\mu\nu}$ is the Ricci's tensor and R is the Ricci scalar. These quantities are all discussed in detail below. G is the Newton's gravitational constant and c is the speed of light.

1.1.1 The Metric

The metric $g_{\mu\nu}$ gives the notion of distance on the manifold, such as:

$$(1.2) \quad ds^2 = g_{\alpha\beta} dx^\alpha dx^\beta,$$

where ds is an infinitesimal distance expressed in the Cartesian coordinate system (t, x, y, z) . As the distance between a points A and B is the same as the distance between B and A , we impose the metric to be symmetric with respect to its indices, for example $g_{xy} = g_{yx}$.

In cases where $g_{\alpha\beta}$ is diagonal with elements 1, the metric is said to be flat and unambiguously described as an Euclidean space. We can verify that the definition of distance satisfies Pythagoras' theorem, which is simply the 2-dimensional case in Euclidean geometry,

$$(1.3) \quad ds^2 = dx^2 + dy^2 = \sum_{\alpha,\beta} g_{\alpha\beta} dl^\alpha dl^\beta,$$

but only if $g_{\alpha\beta}$ is a two-dimensional tensor of the form:

$$(1.4) \quad g_{\alpha\beta} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

The manifold described by the metric has a notion of smoothness (e.g. a cube with edges and corners is different than a sphere) and a notion of angles and distances (e.g. a sphere with constant curvature is different than an ellipsoid) that can be used to describe the geometry and curvature of the universe.

To introduce the notion of true causality (e.g. that past events affect the present), we allow the metric $g_{\alpha\beta}$ to be negative by having one dimension with a sign opposite to the others, e.g.

$$(1.5) \quad g_{\alpha\beta} = \eta_{\alpha\beta} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

We can then make the distinction between points separated in time (time-like, if ds^2 is negative) and points separated in space (space-like, if ds^2 is positive), or moving at the speed of light (null, if $ds^2 = 0$).

1.1.2 Ricci Curvature Tensor

The Ricci curvature tensor $R_{\alpha\beta}$ defines the deviation of the curved spacetime from the flat Euclidean space in terms of volume. To construct it, we start by taking the Riemann curvature tensor $R^\mu_{\nu\mu\sigma}$, measuring the differences between the curvature of our manifold and the traditional Euclidean space [74]:

$$(1.6) \quad R(u, v)w = \nabla_u \nabla_v w - \nabla_v \nabla_u w = \nabla_{[u, v]} w.$$

If there are no differences between the curvature of the manifold and the classic Euclidean space, the Riemann tensor vanishes everywhere (e.g. $R^\mu_{\nu\mu\sigma} = 0$) and the manifold itself is Euclidean (e.g. flat). To extract information about the change in volume of a region of space, we sum $R_{\nu\sigma}$ over all possible index values of its indices, also named 'contracting' in literature, and multiply it by the inverse metric:

$$(1.7) \quad \sum g^{\mu\nu} g_{\nu\rho} = \delta^\mu_\rho = \begin{cases} 1, & \text{if } \nu = \rho, \\ 0, & \text{if } \nu \neq \rho., \end{cases}$$

from which we can build the Ricci tensor:

$$(1.8) \quad R_{\nu\sigma} := \sum R_{\nu\mu\sigma}^{\mu}.$$

1.1.3 Stress-Energy

We relate matter to the geometry of our manifold with the stress-energy tensor $T_{\mu\nu}$, representing the amounts of energy, momentum, pressure, and stress in the space, i.e.

$$(1.9) \quad T_{\mu\nu} = \begin{pmatrix} u & \rho_x & \rho_y & \rho_z \\ \rho_x & P_{xx} & \sigma_{xy} & \sigma_{xz} \\ \rho_y & \sigma_{yx} & P_{yy} & \sigma_{yz} \\ \rho_z & \sigma_{zx} & \rho_{zy} & P_{zz} \end{pmatrix}$$

where u is the energy density, ρ are momentum densities, P are pressures, σ are shear stresses, and the component $T_{\mu\nu}$ shows how much momentum in the μ direction flows in the ν direction through a given point in spacetime.

1.2 Gravitational Waves

Following the definition of the flat space (Eq. 1.5), we can then imagine a “small” perturbation $h_{\alpha\beta}$ as a deviation from the curvature tensor $g_{\alpha\beta}$ from the Minkowski metric η . Explicitly:

$$(1.10) \quad g_{\alpha\beta} = \eta_{\alpha\beta} + h_{\alpha\beta}.$$

Substituting $g_{\alpha\beta}$ in Eq. 1.10 yields the expression for the linearised Ricci tensor:

$$(1.11) \quad R_{\mu\nu} = \frac{1}{2}(h_{\mu,\nu\alpha}^{\alpha} + h_{\mu,\nu\alpha}^{\alpha} - h_{\mu,\nu\alpha}^{\alpha} - h_{\mu,\nu\alpha}^{\alpha}).$$

It is convenient to define the trace-reversed perturbation $\bar{h}_{\mu\nu}$ such as

$$(1.12) \quad \bar{h}_{\mu\nu} = h_{\mu\nu} - \frac{1}{2}h\eta_{\mu\nu}.$$

Evaluating Eq. 1.12 shows that

$$(1.13) \quad \bar{h} = \eta^{\mu\nu}\bar{h}_{\mu\nu} = \eta^{\mu\nu}h_{\mu\nu} - \frac{1}{2}h\eta^{\mu\nu}\eta_{\mu\nu} = h - 2h = -h,$$

which gives this notation the name of ‘trace-reversed perturbation’.

By switching to a Lorenz gauge, it can be shown that Eq 1.2 reduces to 6 components such that

$$(1.14) \quad \square\bar{h}_{\mu\nu} = \frac{16\pi G}{c^4}T_{\mu\nu},$$

where $\square = g_{\mu\nu}\frac{\partial}{\partial x^{\mu}}\frac{\partial}{\partial x^{\nu}}$.

In a vacuum, the stress-energy tensor $T_{\mu\nu}$ vanishes and

$$(1.15) \quad \square\bar{h}_{\mu\nu} = 0$$

for which a plane monochromatic wave propagating in z with speed c is a solution :

$$(1.16) \quad \bar{h}_{\mu\nu} = A_{\mu\nu} e^{i\kappa(ct-z)},$$

with $\kappa = \frac{2\pi}{\lambda} = \frac{2\pi f}{c}$ and f is the frequency of the gravitational perturbation.

The tensor $A_{\mu\nu}$ is symmetric and defined by the wave amplitude.

Following a transverse-traceless gauge transformation on the symmetric tensor $A_{\mu\nu}$, Eq. 1.5 reduces to two independent non-zero components h_{\times} and h_{+} :

$$(1.17) \quad A_{\mu\nu} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & h_{+} & h_{\times} & 0 \\ 0 & h_{\times} & -h_{+} & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

and the gravitational wave can be plus-polarised (if $h_{\times} = 0$), cross-polarised (if $h_{+} = 0$), circularly polarised ($h_{+} = \pm ih$) or elliptically polarised (in all the other cases).

1.2.1 Properties of Gravitational Waves

Akin to electromagnetic radiation, GWs carry energy and momentum (both angular and linear) from the radiation source at the speed of light c and have two independent polarisations (plus and cross). The GW amplitude is characterised by the single dimensionless quantity, h , indicating a fractional change in length of the object that the GW passes through.

Gravitational waves are produced by a non-uniform acceleration of mass. For slowly moving sources ($v \ll c$) in weak gravitational field, we have

$$(1.18) \quad \bar{h}^{ij} = \frac{2G}{rc^4} \int_V \rho(x^i x^j - \frac{1}{3} \delta^{ij} r^2) d^3x,$$

where ρ is the density of the volume d^3x at position x^i , V is the volume of the source and the integral is evaluated at the “retarded time” ($t - \frac{r}{c}$). Eq 1.18 is known as the quadrupole formula for gravitational radiation.

The luminosity of a GW can be approximated by

$$(1.19) \quad L_{GW} \sim \frac{c^5}{G} \left(\frac{R_{Sch}}{R}\right)^2 \left(\frac{v}{c}\right)^6,$$

and shows that BBH mergers, which have dimension of the order of the Schwarzschild radius R_{Sch} , are the most luminous GW sources.

GW amplitudes fall at a rate linearly proportional to the distance they travelled, i.e.

$$(1.20) \quad h \sim \frac{G}{c^4} \frac{\eta E_{kin}}{r},$$

where η measures the asymmetry of the source (and thus the time-varying quadrupole moment) and E_{kin} is the kinetic energy of the fraction of the source producing GWs.

The effect of a gravitational wave (Figure 1.1) is akin to a stretching and squeezing effect, offset by 45 deg.

1.3 LIGO Detectors

This section gives a brief summary of the basic principles of the aLIGO detectors and their main noise sources. A general overview can be found in [92] and more

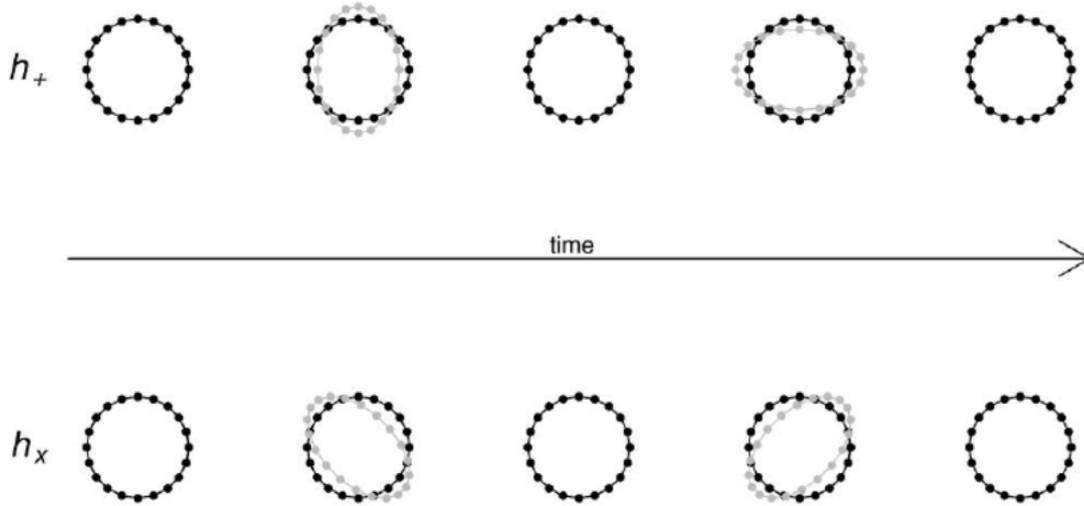


Figure 1.1: Illustration of the effect of a GW passing along the z direction through a ring of particles on the xy plane. The wave's plus polarisation h_+ and cross polarisation h_\times cause an alternating squeezing and stretching of the test masses, offset by 45 degrees. Illustration taken from [27]

details are available in [13]. The Advanced LIGO detectors are four kilometre-long Michelson-type interferometers, where a light beam is split into two perpendicular arms then reflected back and re-combined (Figure 1.2).

The detectors act as a transducer converting GW perturbations into a measurable differential displacement along the arms, proportional to the gravitational wave strain amplitude. As Michelson interferometers do not have a good directional sensitivity to GWs, the strain amplitude is largest for sources with an orbital plane parallel to the detector's plane and located straight above or below. The antenna pattern covers approximately half the sky. The light travel time between the two identical detectors in Livingston, LA, US and Hanford, WA, US is 10 ms. Two pairs of test masses free from displacement noises act as coordinate reference points.

The instruments are calibrated by applying a known force from a laser to the

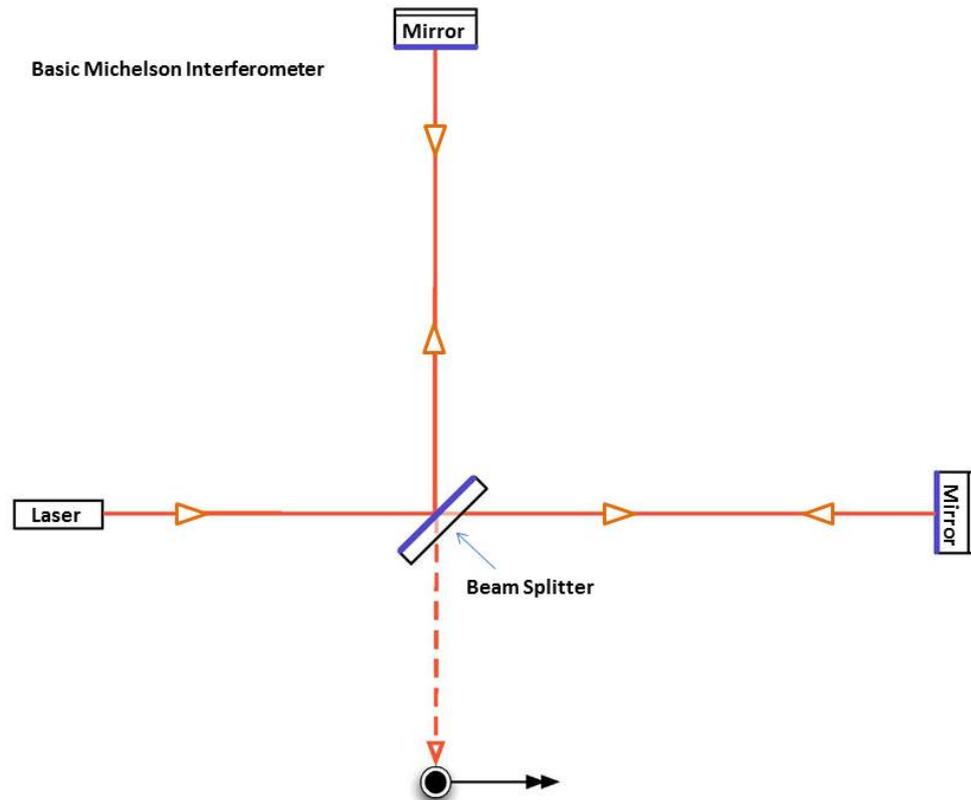


Figure 1.2: Simplified illustration of one of the LIGO detectors. A laser beam is split into two components, then travels to mirrors in perpendicular arms. In cases where the two arms are exactly the same length, the re-combined laser will have destructive interference at the photodetector, and therefore cancel out.

test mass [70] and tracking the effect on an auxiliary optical readout. The system is also used to physically move the mirrors in the same way typical GWs would in a process of ‘hardware injections’. Hardware injections provide critical information on the detection process, allowing consistency checks on the recovery and parameter estimation.

The test masses are suspended by a multi-stage pendulum mounted on actively controlled seismic isolation platforms. The system is built inside ultra-high vacuum chambers. However, classical noise sources are currently not limiting the detector

sensitivity due to the presence of other technical noise sources. [13]

1.3.1 Noise Sources

Figure 1.3 shows the limiting noises as functions of frequency, normalised by the interferometer response.

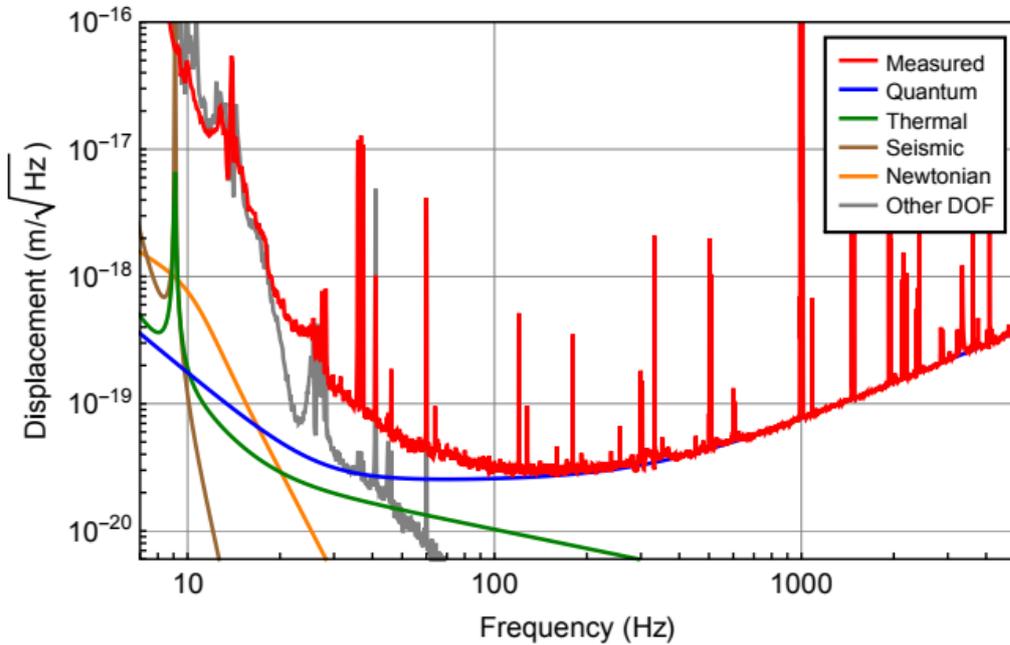


Figure 1.3: Sensitivity of the iLIGO detectors. The quantum noise is the limiting factor at frequencies > 100 Hz. The detector is limited by the other degrees of freedom below 100 Hz, which includes instrumental cross couplings. The strong line features are caused by instruments and calibration lines. Taken from [13]

The detector sensitivity above 100 Hz is limited by photon shot noise, originating from quantum effects.

Radiation pressure noise arises from the photons hitting and transferring their momenta to the mirrors. The quantum radiation pressure noise follows $\frac{1}{mf^2}$, where m is the mass of the mirror and f is the frequency, which is most important at lower

frequencies. Heavier mirrors also help to contain the radiation pressure noise under a frequency threshold.

Counting the photons generated by the laser carries a statistical uncertainty of \sqrt{N} following a Poisson distribution, where N is the number of photons hitting the mirrors at any instant. [92] shows that the shot noise h_{shot} in GW amplitude spectral density is

$$(1.21) \quad h_{shot} = \frac{1}{L} \sqrt{\frac{c\hbar\lambda_l}{4\pi P_{in}}},$$

where L is the length of the interferometer arms, \hbar is the reduced Plank constant, λ_l is the laser wavelength and P_{in} is the laser power.

The shot noise can be minimised in three ways; by increasing the length, L , of the arm and the laser power, P_{in} , and by lowering the laser wavelength, λ_l . Increasing the length of the arms by physically building longer structures is impractical for many reasons. Instead, the LIGO detectors contain optical resonators reflecting the laser back-and-forth along the arm while optimising the effective antenna length, effectively increasing the length of the arm by a factor of ~ 300 . Ever increasing the laser power is similarly unrealistic. Instead, the interferometer is held near a dark fringe and most of the light is reflected back to the laser source by a partially transmissive mirror at the input, leading to an effective power gain of $\sim 40x$. Together, both techniques increase the 20W laser power to an effective 100 kW circulating in each arm.

At frequencies below 100 Hz, the detector at design sensitivity is mainly limited by thermal noise arising due to brownian motion. Thermal noise cannot be completely removed, but by improving the mirror quality Q , its amplitude can be

concentrated in narrow frequency bands around the resonance frequencies ω_0 of the optic- and suspension- vibrational modes

On top of the physical noise, the detectors exhibit non-stationarity manifested as variations in the level and shape of these sensitivity curves over time [12], see Figure 1.4. The non-Gaussian noise transients are caused by complex interactions between instruments and their environment, and can mimic genuine signals when random noise fluctuations manifest simultaneously yet independently in each detector. Transient noise is one of the main factors limiting the sensitivity of gravitational wave searches. This is discussed in detail in Chapter 3.

1.4 Astrophysical Sources of Long-Duration GW

Transients

The recent GW detections directly confirm that GWs could be emitted from binary black hole systems, and could peak in the LIGO frequency sensitivity band during the merging phase. BBH systems are part of a subset of GW sources that are predictable and accurately modelled. GW waveform simulations of such systems can be generated by varying the most sensitive parameters like system orientation, eccentricity, individual masses or spins. By comparing the detector output with many iterations of different modelled systems, BBHs can be detected more reliably than unmodelled systems. GWs emitted by such systems are also short-lived inside the LIGO sensitivity band.

In contrast, astrophysical events that could potentially generate minute-long GW emissions are not as well understood and cannot be modelled accurately, hence

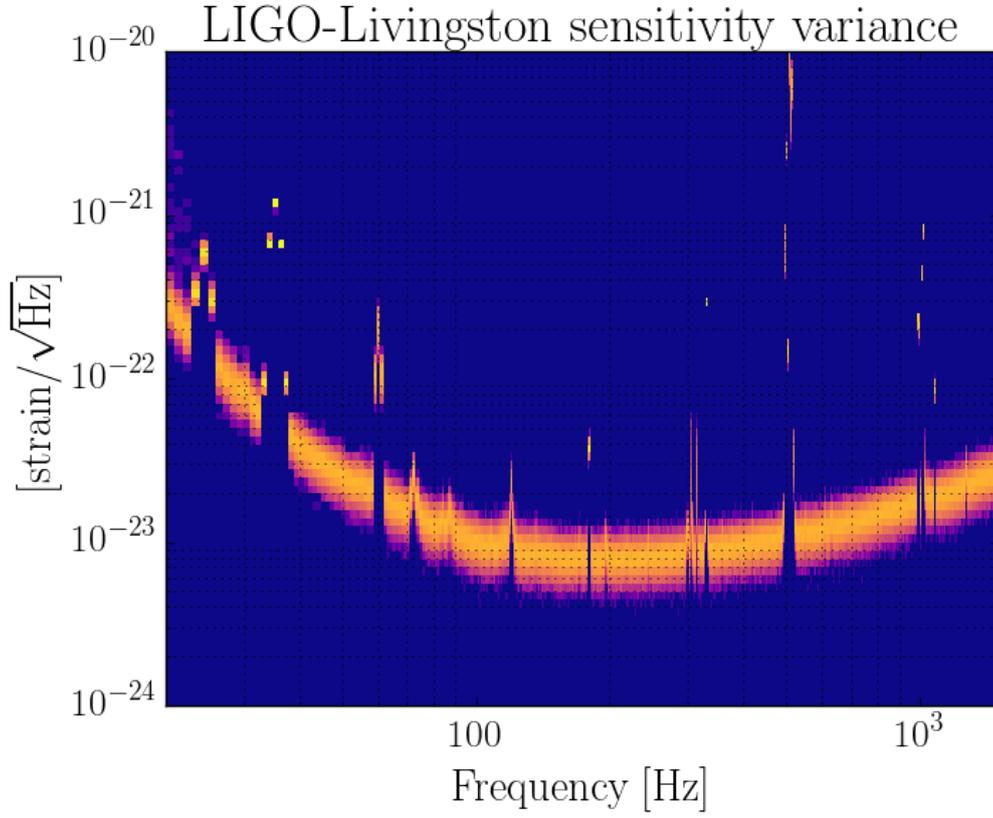


Figure 1.4: : Spectral variance of the advanced LIGO detectors during ~ 1 hour around GW150914 for the Livingston detector (L1). The sensitivity of the detector typically fluctuates by a factor of 5, with the more noisy frequency bands exhibiting variance of few order of magnitude (see the 500 Hz band for example.)

speculations have to be made to allow for numerical solutions. However, Burst-type models are routinely used to characterise the efficiency of burst searches as even inaccurate model signals used in this context have properties consistent with GW emissions such as duration, energy profile or frequency evolution that can be used to give insights into search pipelines.

Astrophysical sources of long GW transients are often associated with complex dynamics and hydrodynamic instabilities in the context of core collapse supernovae

and long gamma-ray bursts (GRBs) [81, 99].

Turbulences in the accretion disk of a black hole could lead to GW emissions [101]. Through strong magnetic fields, the accretion disk around a Kerr black hole forms clumps of matter that are then separated by GW emission spinning down the black hole [102]. GWs emitted following this model have anti-chirp like characteristics, with exact form depending on parameters such as the mass of the black hole, its spin and the fraction of the accretion disk that forms clumps. As the strong coupling and the amplitude of the associated GW emission are described as optimistic [91], GWs from realistic accretion disk instabilities will likely diverge from the model by several orders of magnitude in amplitude and have a wider bandwidth. [103] indicates a detectability of the model to ~ 100 Mpc with aLIGO design sensitivity. As long-duration GRBs have an estimated rate of about $300 \text{ Gpc}^{-3}\text{Yr}^{-1}$ [84], such GW signals could be observed within a few years.

Magnetic deformation of fast spinning neutron stars could generate GW emissions with anti-chirp behaviour [57]. The exact form depends on the spin frequency of the neutron star and its magnetic ellipticity. Newborn spinning neutron stars in unstable configurations could also give rise to non-axisymmetric deformations producing minute-long GWs [85]. Due to their weaker amplitude, such GW signals could be detected with aLIGO sensitivity only if originating from our galaxy.

GRAVITATIONAL WAVE DETECTION BY SPHERICAL RADIOMETRY

2.1 Introduction

Previous gravitational-wave searches with the LIGO detectors have targeted transient signals with durations less than a few seconds [7, 9], and persistent signals with durations up to hour-long and more [8, 10]. This leaves the regime of transient GW signals with durations of 10-1000 seconds relatively unexplored [1, 56, 50].

Searches that are used to detect recent GW events ([13, 14, 26]) make use of pre-computed banks of simulated signals ¹ following the better-understood GW signatures of these astrophysical events [66]. They then compare the observed data from the LIGO detectors with the models to find potential matches (for example, [15]). The technique is extremely sensitive to signals matching the computed model

¹We note that short duration transient GW search pipelines also saw the events

templates, but virtually insensitive to GW events that cannot be confidently modelled [62], or are simply not part of the pre-computed banks.

In contrast, the physics behind potential sources of minute-long GWs is not well understood, and as such no accurate GW emission models are available. Searches for unmodelled GW transients need to use alternative methods to look more generally for a localised excess of power [30], and must rely on minimal assumptions about the signal. Therefore, they need to use techniques that are not specific to any GW model but encompass all possible signals as generally as possible [20]. They are thus overall less sensitive to modelled GW signals than modelled searches but have a larger reach with broader potential for unexpected discoveries.

Both modelled and unmodelled GW searches operate on either of two bases; when the source is known (from an electromagnetic (EM) counterpart such as GRBs [2] or supernova (SN) [16]), the time and sky position of the event are used to specifically target the GW signal. This gives a relatively small parameter space to analyse, so targeted searches from EM counterparts provide good sensitivity and offer a promising future to detection success. When the target is not known *a-priori* (i.e. without an EM counterpart), searches must look over all-sky and observe continuously, resulting in a much larger parameter space and increased computational cost.

The search pipeline described in this chapter specifically targets unmodelled gravitational-wave bursts with durations in the 10-1000 second range over all-sky, motivated by the strong potential for GW detection in this regime [98]. Using a novel approach based on spherical radiometry, the pipeline is competitively fast and introduces a powerful new signal consistency measure well suited for rejection of background noise [51].

In this chapter, we introduce the formalism needed to take the detector data streams from time domain to Fourier, then into the spherical harmonic domain. We derive quantities to infer the likelihood of a GW signal in the data and introduce principles behind background rejection. We then introduce the working pipeline, X-SphRad, and detail the analysis steps and the detection procedure.

2.2 Gravitational Wave Analysis

The most straightforward approach to GW analysis with a network of detectors is to generate a list of candidate events in each detector, then flag events happening simultaneously with overlapping properties in all detectors [30, 96]. The properties of GW candidates can then be extrapolated by a weighted average of the individual properties in each detector. This method is fast but sensitive to glitches that are simultaneous in individual detectors [9].

Another method is to combine the data from the individual detectors prior to generating the list of candidate events [39, 60], then measure the similarity between detectors in the network in individual frequency- and time- regions. Glitches are rejected by a measure of coherence [95]. The detectors are naturally weighted by their relative sensitivity which simplify further computations, but the procedure is slower as the data needs to be correlated first. As GWs travel at the speed of light, the difference in time of the detection of an event in the detectors allows for triangulation of the source in the sky. The proper correlation of two or more GW detectors depends on measuring the time difference, and thus the relative sky position of the emitting source(s). The estimated sky position is also used to enable the identification of a GW, and subsequent rejection of background noise.

To account for sky position of the GW sources and compute the correct correlation, typical all-sky GW searches map a grid of points in the sky and analyse the data coherently as coming from each point. Because there are as many coherent combinations as there are resolvable positions on the sky for the network, the method is computationally costly.

The pipeline described in this chapter drastically lowers the computational time of all-sky searches by removing the sky position dependency from the correlation of the detector outputs in the spherical harmonic domain [40]. The spherical radiometry approach takes advantage of the fact that sky maps in gravitational-wave searches show strong correlations over large angular scales in a pattern determined by the network geometry [50]. Computing sky maps indirectly through their spherical harmonics minimises the redundant calculations.

The pipeline is based on X-PIPELINE [95]; a MATLAB-based package that contains a suite of pre- and post-processing techniques that can be applied on event triggers that are output by the novel Spherical Radiometer (SphRad) engine, developed by [50]. X-PIPELINE is standard in the field and widely used by the community in SN searches [58], GRB searches [21], fast radio-burst triggered searches [17], and is being used for the all-sky O1 search [22]. X-SphRad underwent internal review within the LIGO collaboration.

2.3 Formalism

This section is adapted from [95] and [50].

For a GW with components h_+, h_\times coming from a direction Ω and detected by a detector α with antenna response $F_\alpha(\Omega)$, the data output is be the linear combination:

$$(2.1) \quad d_\alpha(t + \delta t_\alpha(\Omega)) = F_\alpha^+(\Omega)h_+(t) + F_\alpha^\times(\Omega)h_\times(t) + n_\alpha(t + \delta t_\alpha(\Omega)),$$

where n_α is the noise contribution to the output and $\delta t_\alpha(\Omega)$ is the time delay due to the speed of GW propagation between an arbitrary reference.

For convenience, we can introduce the nose-weighted quantities in Fourier space, i.e.

$$(2.2) \quad \tilde{d} = \tilde{d}_{w\alpha}[k] = \frac{\tilde{d}_\alpha[k]}{\sqrt{\frac{N}{2}S_\alpha[k]}},$$

$$(2.3) \quad \tilde{n} = \tilde{n}_{w\alpha}[k] = \frac{\tilde{n}_\alpha[k]}{\sqrt{\frac{N}{2}S_\alpha[k]}},$$

$$(2.4) \quad F = \tilde{F}_{w\alpha}^{+,\times}(\Omega, k) = \frac{\tilde{F}_\alpha^{+,\times}(\Omega)}{\sqrt{\frac{N}{2}S_\alpha[k]}},$$

where we have used

$$(2.5) \quad x(t) \rightarrow \tilde{x}(F) \equiv \frac{1}{\sqrt{2\pi}} \int dt e^{-2\pi i f t} x(t) dt.$$

To simplify the notation, we can rewrite Eq. 2.1 in the frequency domain:

$$(2.6) \quad \tilde{d} = F\tilde{h} + \tilde{n}.$$

We can now introduce the probability $P(\tilde{d}|\tilde{h})$ of obtaining the whitened data \tilde{d} in a single time-frequency pixel in the presence of a GW h from known sky direction:

$$(2.7) \quad P(\tilde{\mathbf{d}}|\tilde{\mathbf{h}}) = \frac{1}{(2\pi)^{D/2}} e^{-\frac{1}{2}|\tilde{\mathbf{d}}-F\tilde{\mathbf{h}}|^2}.$$

The likelihood ratio L is then the probability of obtaining the whitened data under the null hypothesis, where we treat an unknown GW h as a free parameter to be fit to the data. We find that the value \hat{h} that maximises L is

$$(2.8) \quad \frac{\partial L}{\partial h} |_{h=h_{max}} = 0 \Rightarrow h_{max} = (F^T F)^{-1} F^T d,$$

or the Moore-Penrose pseudo-inverse [69].

With,

$$(2.9) \quad L = \frac{P(\tilde{\mathbf{d}}|\tilde{\mathbf{h}})}{P(\tilde{\mathbf{d}}|\mathbf{0})} = \frac{\frac{1}{(2\pi)^{D/2}} e^{-\frac{1}{2}|\tilde{\mathbf{d}}-F\tilde{\mathbf{h}}|^2}}{\frac{1}{(2\pi)^{D/2}} e^{-\frac{1}{2}|\tilde{\mathbf{d}}|^2}} = e^{\left[\frac{1}{2}|\tilde{\mathbf{d}}|^2 - \frac{1}{2}|\tilde{\mathbf{d}}-F\tilde{\mathbf{h}}|^2\right]}.$$

Substituting this into Eq. 2.8 gives the max or “standard likelihood”

$$(2.10) \quad E_{SL} = 2 \ln L(\tilde{h}_{max}) = |\tilde{\mathbf{d}}|^2 - |\tilde{\mathbf{d}} - \tilde{F}\tilde{h}_{max}|^2 = 2\tilde{\mathbf{d}}F\tilde{h}_{max} - |F\tilde{h}_{max}|^2 = 2\tilde{\mathbf{d}}F(F^T F)^{-1}F^T d,$$

where $F(F^T F)^{-1}F^T$ acts as a projection of the data. For the case of 2 non-aligned detectors, it can be shown that the max-like solution (Eq. 2.8) is equal to the total energy in the data, or

$$(2.11) \quad E_{tot} = \sum_k |\tilde{\mathbf{d}}|^2,$$

and is uniquely a sum of the autocorrelation terms.

However for more general cases in a network of 3+ detectors, Eq. 2.8 contains both auto- and cross- correlation terms and thus need to be computed over the whole sky. These computations can be fast and efficient if we move into the spherical harmonic domain.

We can write the auto- and cross- correlation terms ξ_{ij} from the instruments i and j at positions \vec{r}_i and \vec{r}_j and their output data streams g_i and g_j delayed by the sky position \hat{s} , over an interval T centred at time t

$$(2.12) \quad \xi_{ij}(\hat{s}) = \int_{t-T/2}^{t+T/2} g_i(t' - \vec{r}_i \cdot \hat{s}) \ g_j(t' - \vec{r}_j \cdot \hat{s}) \ dt'.$$

The geometric delay is a linear transformation. We can separate out the delay T in data streams due to a source with a sky position as :

$$(2.13) \quad \vec{g}(\hat{s}) = T(\vec{r} \cdot \hat{s}) \cdot \vec{g}$$

and then rewrite Eq. 2.12 as :

$$(2.14) \quad \xi_{12}(\hat{s}) = \frac{1}{N} \vec{g}_1^t \cdot T^T(\vec{r}_1 \cdot \hat{s}) \cdot \vec{g}_2 \cdot T(\vec{r}_2 \cdot \hat{s}),$$

where the matrix $T(\vec{r} \cdot \hat{s})$ holds the delays in the time-series due to the sky position. Parseval's theorem allows us to relate the power in different domains; i.e. the total energy for all times in the time domain is equal to the total energy for all frequencies in the frequency domain:

$$(2.15) \quad \sum_{n=0}^{N-1} |x[n]|^2 = \sum_{k=0}^{K-1} |x[k]|^2$$

Then Eq. 2.14 can be written in its discrete form as

$$(2.16) \quad \xi_{12}(\hat{s}) = \frac{1}{N} \sum_{j=0}^{n-1} g_1[\mathcal{J}; \hat{s}] g_2[\mathcal{J}; \hat{s}].$$

Since Eq. 2.16 is a function of position \hat{s} on the sphere, it can be approximated by a sum over spherical harmonics $Y_{(lm)}\hat{s}$ up to some order l_{max} and Eq. 2.16 can be approximated by

$$(2.17) \quad \xi_{ij}[q]^{(lm)} = \frac{1}{N} \sum_{l=0}^{l_{max}} \sum_{m=l}^l \sum_q \left[\tilde{g}_i[q] \tilde{g}_j^*[q] \left(\frac{1}{N} \tilde{T}_q(\vec{r}_i \cdot \hat{s}) \tilde{T}_q^*(\vec{r}_j \cdot \hat{s}) \right)^{(lm)} \right] Y_{lm}(\hat{s}),$$

where l and m are the indices of the spherical harmonic Y_m^l , and k is the Fourier transform of T that applies the appropriate time delay for sky position \hat{s} . The Fourier Transform makes the convolution $g \cdot T$ a simple multiplication in the frequency domain. Each frequency bin q contains information about the time, frequency and direction decomposition of the detector output. Dependency to sky location T is isolated and we can Fourier transform the data without relationship to the sky position. The details of the time-frequency analysis in X-SphRad are discussed in Section 2.4.3. Here, we note that we use the following ad-hoc statistic for selecting candidate events in the time-frequency map. For example, the spectrogram for a network of two detectors would be constructed as each pixels having a value of:

$$(2.18) \quad N = \xi_{1,2}^{(lm)} + \xi_{1,1}^{(lm)} + \xi_{2,2}^{(lm)}.$$

2.4 Pipeline Workflow

The spherical radiometry routines are integrated into X-PIPELINE, which handles the pre- and post-processing of data. The pre-processing step consists of gathering data and segmenting it into more manageable, shorter chunks, as well as pre-conditioning the data with a whitening filter. The X-SphRad core does the spherical harmonic (SH) decomposition and generates time-frequency maps, computes statistics about pixel energies then identifies clusters of interest. The post-processing step optimises cuts and reports sensitivity and events of interest (Fig. 2.1). The pipeline is set up to analyse a network of 2 or 3 detectors over all-sky. It is also analysing at all-time, so we are interested in all abnormal events in the analysed time.

2.4.1 Data Pre-Conditioning

The first set of tasks for the pipeline is to list the data available and segment it into more manageable smaller-duration blocks, then whiten the block and Fourier transform it.

Figure 2.18 shows that X-SphRad measures the coherence between detectors, and thus needs coincident data for all the detectors in the network. In cases where one or more datastreams are not fully available during the time block, the block is entirely skipped. We do not have a measure of the coherence and therefore the data segment must be skipped, even for only one instant of single detector downtime. Even a minimal loss of detector uptime can significantly impact the analysis if happening at intervals around the block length.

A zero-phase linear predictor filter whitens the data and estimates the power spectrum to remove standing “power lines” due to instrumental interference [4]

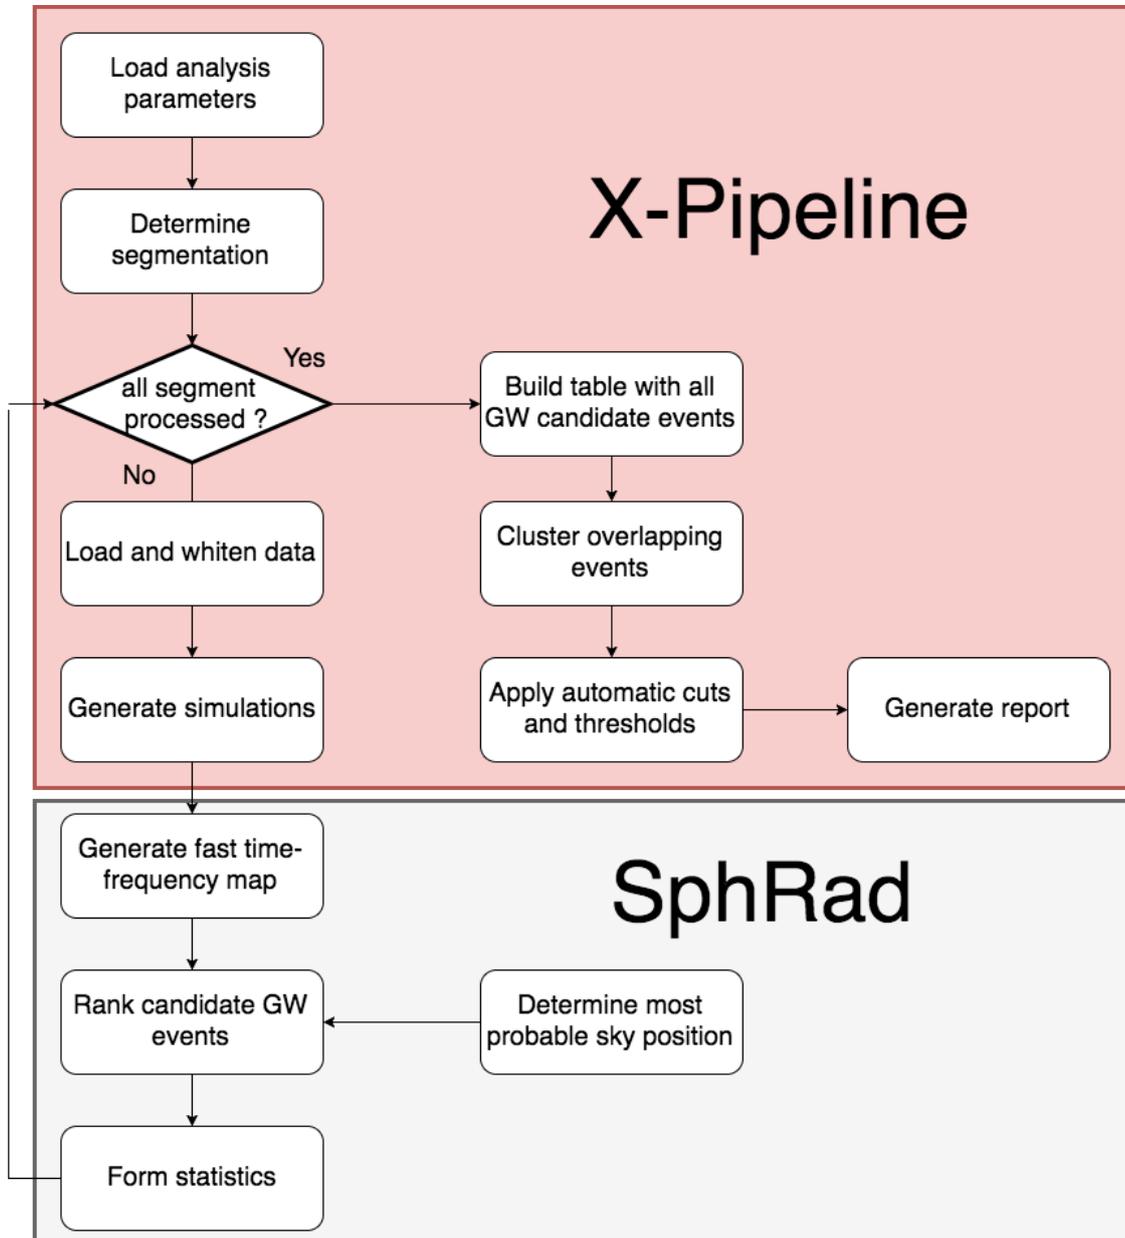


Figure 2.1: Workflow of the X-SphRad pipeline. X-pipeline is used for the pre- and post- processing, while the spherical radiometer SphRad is use as core detection engine.

without inducing the nonzero phase response of classic linear predictor error filters [43]. We use a 2 second filter time, and discard 8 seconds at each end of the block length, and so consecutive data blocks are overlapped by multiples of 16 seconds to ensure no signals are missed. This step also minimises data loss due to single detector uptime loss. The filter estimates the power spectrum and essentially normalises energies at each frequency over the block time. During the process, the data is also resampled and band-passed with minimal phase shifting. The whitening is done in the frequency domain, then transformed back in the time domain.

The optimal duration of the block depends on the targeted events and should thus be chosen carefully. Short duration blocks are individually quicker to process, although Matlab startup times and in/out reading reduces the amount of time saved. However, each block needs to be large enough so that events of interest are fully contained within a block, yet longer blocks risk the power spectrum estimation (from the whitening filter) to be contaminated by the non-stationary detector noise. For example, supernovae may emit sub-second GW signals [80] making the optimal block length for such observations a few orders of magnitude smaller than for target signals lasting for several minutes. To mitigate the power spectrum estimation error and maximise event recovery, the current pipeline is set up to join up to two blocks together in the later steps. We note that the rotation of the Earth has a negligible effect on searches, as the rotation of about 0.25 deg/min is contained within the pipeline resolution of about 3 degrees.

Finally, the data is Fourier transformed into the frequency domain. As the detector is discretely sampled, it is convenient to use discrete notation. The Fourier transform $x[k]$ of the time series $x[j]$ for N data points can be written as

$$(2.19) \quad \tilde{x}[k] = \sum_{j=0}^{N-1} x[j] e^{-\frac{i2\pi jk}{N}}.$$

The resolution in time-frequency is limited by the following parameter used for δt and δf . Bound by uncertainty principle [79], short time-duration fast Fourier transform (FFT) will lose resolution in frequency domain, while longer time-duration FFT gain resolution in frequency :

$$(2.20) \quad 1 = \delta t \times \delta f.$$

As always when using FFT, the overlap needs to be set appropriately to compensate for the taper window used to prevent power leakage [37]. The time-frequency resolution should be adapted to the target signals to extract the maximum amount of information and reduce noise corruption.

2.4.2 Sky Position Dependency

The delay matrix is constructed for each detector-pair baseline in the frequency and spherical harmonic domains to offset the time series according to the sky location (Eq. 2.17). The input time series corresponding to a pair of detectors are transformed to the frequency domain to give $g_i^*[q]$ and $g_j^*[q]$, then multiplied together at each frequency bin q . Each bin is multiplied by the harmonic coefficients in the corresponding frequency bin of the delay matrix, and yields the harmonic coefficients for the angular distribution of cross power for that frequency bin. Summing across bins generates the baseline direction-dependent integrated cross-power in the form of a single vector of coefficients of its expansion in spherical harmonics [50].

2.4.3 Time-Frequency Maps

The pipeline then generates a spectrogram, or time-frequency map (TFmap), from the spherical harmonic expansion using FFT (Eq. 2.19). For current long-duration GW burst searches, the time-frequency transformation of the data is a crucial step. It is justified, as target GW signals appear as thin tracks of extra power on a TFmap, and hence it is band limited and localised. Separating GW candidates is thus more straightforward, as they can be physically located in different time-frequency regions of the TFmap [43].

The chosen time-frequency resolution determines the aspect ratio of pixels on the TFmap with the relationship $\delta t \delta f = 1$. Longer time integration produces good resolution in the frequency domain but poor resolution in the time domain, while shorter time integration inversely yield poor frequency resolution. Due to the nature of the Fourier transform, each pixel on the time-frequency map indiscriminately accumulates background noise and potential signals within the space $\delta t \times \delta f$. The time-frequency resolution should then be selected to allow for minimum noise pollution in both the time and frequency domains. The time-frequency resolution used in the FFT needs to be optimised for maximum signal-to-noise ratio (SNR) of target events within each time-frequency bin and so depends on the target signal properties.

The value of each pixel P on the time-frequency map is then a sum along the spherical harmonic coefficients c representing the energy in a time-frequency bin $dt \times df$, where

$$(2.21) \quad P_{t,f} = \sum_{i,j=0}^d \xi_{ij}^{(lm)} = \sum_{l=0}^{l_{max}} \sum_{m=-l}^{m=l} (c_{lm})^2.$$

In practice however, tests done in [50] showed that the computational cost of

the full cross- and auto- correlations of data streams with sky & frequency dependant weight factors (Eq 2.10) is significantly more expensive, with an increase in processing time proportional to the length of the block being analysed.

The cross-correlation variation over the sky is mostly due to changes in time-delay between detectors due to sky position, and so the correlation is approximately constant on a ring in the sky of constant delay (Figure 2.2). We choose the coordinate system so that the inter-detector baseline is along the \hat{z} axis, and so the rings of constant delay correspond to the $m = 0$ mode in the spherical harmonic domain. By doing so, we can get an approximate cross-correlation by ignoring the $m > 0$ modes and computing the $m = 0$ only mode for each detector pair [50]. Figure 2.3 compares time-frequency maps generated with the full and reduced harmonic. It shows no difference visually.

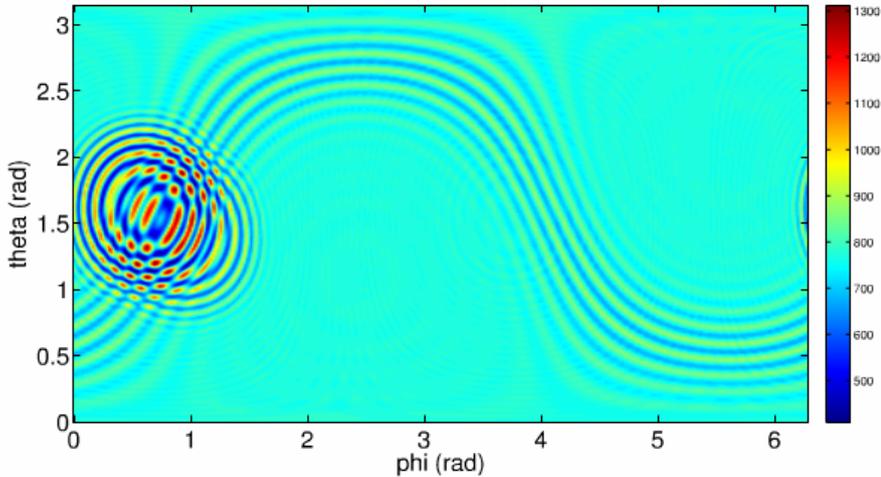


Figure 2.2: Skymap of a network of 3 detectors. Each pair of detectors produce a ring of constant delay where the crossing of the rings is the most probable sky location of the signal. Taken from [50]

To accelerate the time-frequency map generation, both the sky & frequency

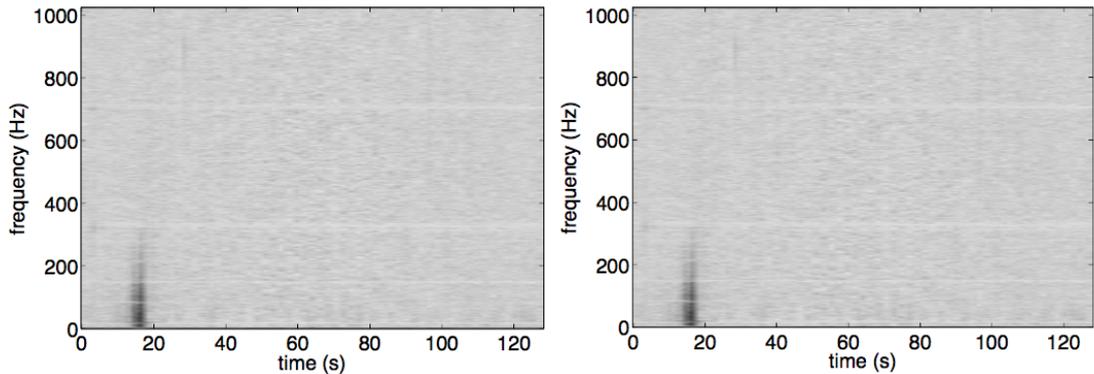


Figure 2.3: Comparison of time-frequency map using the full (Eq. 2.17) (left) and reduced (Eq. 2.18) statistic (right). The two agree well, as no difference is visible between features of both maps. Taken from [50]

weights are initially discarded. The higher order modes $m > 0$ are also discarded, as they were shown to contain relatively little power compared to the mode $m = 0$, and subsequently have little impact on the final pixel value output. The resulting time-frequency map is visually identical to one generated by the full Eq. 2.10, but discards all the spatial information gathered from the spherical harmonic expansion as the step is equivalent to averaging over the whole sky.

2.5 Clustering

Typical all-sky analyses identify candidate GW events by searching for an excess of power [30] on a time-frequency map. The process is generally referred to as *clustering* [96].

X-SphRad flag potential GW candidates by dividing the time-frequency map in 10 time-regions of equal time-frequency area. For each region, a core threshold identifies pixel ‘seeds’, from which branches grow in all possible directions up to reaching a pixel with value lower than the halo threshold. The selected pixels are

then clustered together on the condition that they share either an edge or a corner in the TF plane (‘next-nearest-neighbours clustering’) [31].

The branches found via this method are then ranked by total energy, and the top x are classified as candidate events and selected for further investigation. The double thresholds are robust against slow gaussian-like noise variations in data, but very loud and/or wide-band glitches that cover large time-frequency areas can saturate the clustering algorithm and cause weaker potential GW signals to be ranked lower than the top x louder clusters due to the limit of the number of brightest pixels that are selected for consideration.

Coherent statistics of interest (Eq. 2.17) are then computed using the full (l, m) decomposition for the x clusters with the largest summed values of correlated power P (Eq. 2.18).

2.5.1 Background Estimation

To measure the significance of candidate GW events and differentiate between noise and signal, the pipeline estimates the background using the time-shifted version of the data itself by using the implementation of the time slide technique, generally adopted in searches for GWs [9, 105]. The data stream from one arbitrarily chosen detector in a network is shifted by non-physical time offsets and analysed as real data, as previously described. The time offsets need to be at least larger than the time taken by a GW to travel from one detector to the other to avoid introducing a correlation that could be attributed to the sky position of the GW source.

Because they originate from non-physical realisation of the data, all GW candidates generated can be safely treated as noise fluctuations, and are used to estimate

the non-stationary background noise over the data period. When a full background estimation is needed, for example to detect a GW with a confidence level above 5 sigma, the background analysis is by far the most computationally intensive operation in the whole pipeline and the reason why a key emphasis is put on the speed of X-SphRad.

2.5.2 Software Injections

Simulated signals are commonly used to measure the sensitivity of burst pipelines, giving the distance at which a pipeline can reliably detect a simulated signal drawn from an isotropic distribution of sky locations and orientations. The efficiency at a given distance is defined by the fraction of signals at that distance passing all background rejection cuts, and with a ranking statistic equal to or larger than a value corresponding to the given false alarm rate estimated from the background noise.

The model signals are added to the individual detector independently in the time domain, and evenly spaced time over the analysed data period. The amplitude of the simulated signals is multiplied by a factor corresponding to a distance between the GW source and the LIGO detector logarithmically spaced between 100 kpc to 1000 Mpc for astrophysical models. The distance is meaningless for ad-hoc models, but a similar factor was chosen to modify the amplitude of the GW. The source location (ϕ, θ) on the sky and orientation relative to the earth are drawn from the isotropical distributions:

- $\cos(\theta)$ uniform on $[-1, 1]$
- ϕ uniform on $[0, 2\pi)$

- $\cos(\iota)$ uniform on $[-1, 1]$
- ψ uniform on $[0, \pi)$

2.5.3 Coherent Cuts and Tuning

Working in the spherical harmonic domain offers an interesting byproduct; as GWs are by definition coherent and so have a preferred direction in the sky, their energy is spread over all harmonic coefficients. By contrast, glitches are not necessarily coherent between detectors and so do not have preferred direction in the sky. Therefore, their energies are mainly concentrated in the lowest order mode $l = 0$. We define the incoherent energy I_{sh} as the power contained in the $l = 0$ mode and the coherent energy E_{sh} as the sum of the power contained in all the higher order mode $l > 0$.

We then construct the statistic

$$(2.22) \quad \Gamma = \frac{\sum_{l=1}^{l_{max}} \sum_{m=-l}^l |c_{lm}|}{|c_{00}|},$$

where c_{00} is the energy average over all sky (or I_{sh}) and the numerator is the sky energy (or E_{sh}). Thus the ratio of $\frac{E_{sh}}{I_{sh}}$ is a good indicator for rejecting artificial noise [51]. We can thus make a linear cut on the ratio as

$$(2.23) \quad E_{sh} \geq \kappa I_{sh},$$

where κ is a constant selected during the automated tuning procedure.

However, low energy events have shown a tendency to get corrupted by random noise fluctuations and have slightly more power in the coherent energy E_{sh} than in the incoherent energy I_{sh} . For this reason, we apply an additional coherent cut,

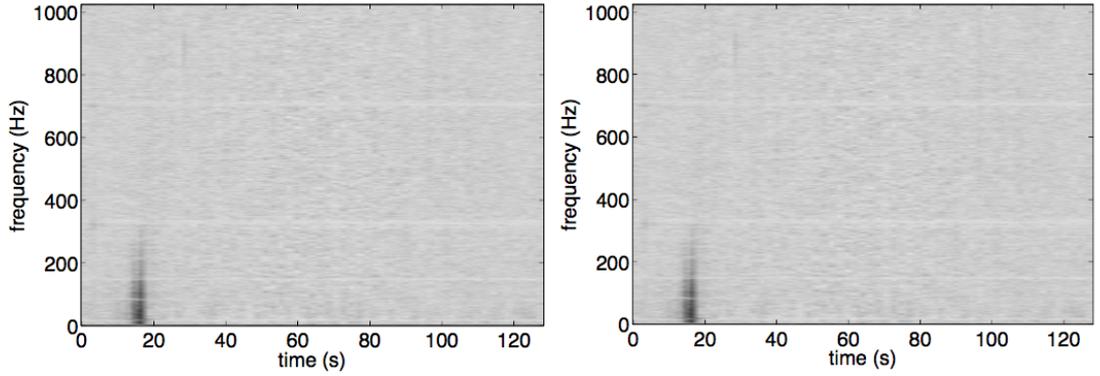


Figure 2.4: Sky map energy E_{sh} versus incoherent energy I_{sh} for simulated data (black) and injections of various amplitudes (coloured). The coloured crosses on top of the background are due to the post-processing misselecting background triggers as injected signals. Taken from [50]

alpha cut, constructed to remove low energy glitches. The form of the cut was chosen empirically, and was shown to be good at removing low energy glitches [97];

$$(2.24) \quad r = \frac{2|E_{sh} - I_{sh}|}{(E_{sh} + I_{sh})^{0.8}} + 1 \geq threshold,$$

where the threshold is determined during the automated tuning procedure.

The sensitivity of the pipeline is measured by the distance at which it can confidently detect a set of test GW waveforms. Such waveforms are modelled approximately, hence they do not suffice in the case of a modelled search but are acceptable for characterisation of unmodelled searches. For a test waveform to be detected, or ‘recovered’, we require at least a few pixels of the event to have a significance larger than the significance of loudest background noise.

There is a trade-off in tuning the thresholds for searches; a lower threshold means accepting more background noise, whereas higher thresholds allow the possibility of rejecting potential signals in the data. The pipeline empirically tests a set of

CHAPTER 2. GRAVITATIONAL WAVE DETECTION BY SPHERICAL RADIOMETRY

values for each threshold on our statistic Z and outputs the threshold that yields the maximised amount of signal recovery on average versus accepting background noise events generated in the background estimation step.

O1 LONG-DURATION BURST SEARCH

3.1 Overview

We apply the X-SphRad pipeline to the first advanced LIGO (aLIGO) science run in the context of an all-sky search for unmodelled transient GWs of long duration. We present the results of the search using the LIGO Hanford - Livingston second generation detector network spanning from September 2015 to January 2016. The parameter space considered by the search covers frequencies from 24 Hz to 1000 Hz, with target gravitational wave transients lasting from a few seconds up to ~ 1000 s.

This search is characterised by making minimal assumptions about the signal. Basic hypotheses following the physical properties of GWs made in Chapter 1 (e.g., that GWs have two polarisations that propagate at the speed of light), intrinsic parameters such as waveform or extrinsic parameters such as time of occurrence or propagation direction have not been targeted specifically. This is in contrast to

searches for compact binary coalescence that manage to be less sensitive to random noise fluctuations by targeting specific waveform morphologies (see for example [100] and [87]).

The search reports a few low-significance candidate events that are consistent with background noise. As a consequence, no significant events were observed and we use the loudest event procedure in order to set upper limits on the rate of long duration GW transient signals at 90% confidence.

We start by describing the data set and the data selection process, then expose data quality issues. We show examples of non-stationary background noise present in the data and detail steps taken to minimise their impact on the search. We then detail X-SphRad's handling of the search parameters and the candidate event selection building on Chapter 2. We detail a procedure for estimating the search background due to the coincidence of non-stationary noise in the detectors, using non-physical time shifts. We also determine the sensitivity of the search with a isotropic distribution of simulated GW bursts, both astrophysically motivated and *ad-hoc*. Finally, we report the upper limits from the search and compare our sensitivity to the other long duration pipelines used in aLIGO and show X-SphRad to be competitive.

3.2 Observing Run 1

The data were taken from the first aLIGO Observing run (henceforth referred to as O1), which took place from September 12, 2015 to January 19, 2016. The O1 run consisted of 130 days open for joint data collection between the two-detector network at Hanford, WA (H1) and Livingston, LA (L1) with the highest sensitivity to date [13]. The O1 run has been marked by the first direct detection of the gravitational

wave events GW150914 [13], followed by GW151226 [14] and the candidate event LVT151012 [18].

3.2.1 Analysed Science Segments

The LIGO detectors need to achieve a *cavity lock* state ¹ to be sensitive to a GW. A multitude of factors hinder the detector’s ability to reach or maintain this state, reducing the actual observing time for which data can be processed and interpreted with confidence. In this text, we follow the adopted convention and use the term *science mode* to refer to periods of time when a detector is gathering data with scientific intent. Over the course of the O1 period, H1 observed in science mode 62% of the time, while L1 observed 55% of the time (Figure 3.1).

The difference in science mode operation between the two detectors is explained in part by the difference in geographical location, and thus naturally occurring weather events. For example, the principal causes for lost science mode time in H1 are locking issues, accounting for 18%, and environmental disturbances, accounting for less than 10%. In contrast, the main causes at L1 are environmental, accounting for nearly 20%, and locking issues, accounting for 16%. At both sites, the deliberate downtime for maintenance work accounts for less than 5% of the run [93].

The total time over which both detectors were coincidentally in science mode is 49 days. As these periods are not continuous, they will be referred to as coincident science segments to avoid confusion with the O1 run total observing time. Table 3.1 shows the O1 live times. The final duration of the data analysed by X-SphRad totals 49.1 days, or 38% of the time period covering O1. We note that the total time processed by X-SphRad is smaller than the coincident network time, as the pipeline

¹In the cavity lock state, the length of the LIGO arms is kept precisely constant [13]

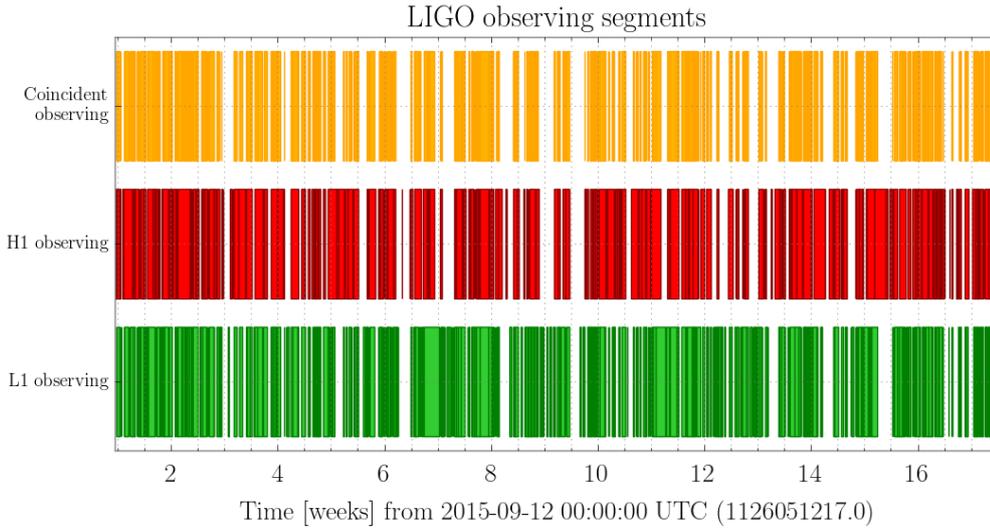


Figure 3.1: Representation of aLIGO observing time during O1 for both the H1 and L1 detectors, as well as coincident observing time from September 12, 2015 to January 19, 2016. H1 observed in science mode during 78.8 days (62% duty cycle) and L1 during 67.2 days (55% duty cycle). The total network duty cycle for science mode is 40%. The pictured segments do not take into account subsequent data quality vetoes. See also Table 3.1. Generated with [76]

require a minimum segment length of 512s of continuous coincident science mode data for processing (Section 2.4.1).

3.2.2 Sensitivity

The first advanced LIGO run has shown an amplitude sensitivity increase by a factor of approximately 3 or 4 compared to initial LIGO runs [4]. Figure 3.2 and Figure 3.3 show representative noise amplitude spectra for H1 and L1 during O1. The strong narrowband features in the spectra are caused by resonances in the instruments such as the violin modes of the suspension wires, the roll and bounce modes of the suspensions, the AC power line and its harmonics, and the calibration lines [13].

To put the instrumental sensitivity in an astrophysical context, the angle-

Table 3.1: Breakdown of dataset total duration and science mode time after applying data quality requirements and coincidence. The time processed by X-SphRad is slightly less than the H1-L1 coincident time due to the pipeline requiring a minimum of 512 s of continuous data for analysis.

O1 period	129.7 days
H1 science mode	78.8 days
H1 (science - CAT1)	76.8 days
L1 science mode	67.2 days
L1 (science - CAT1)	65.8 days
H1-L1 coincident time	51.3 days
X-Sphrad processed time	49.1 days

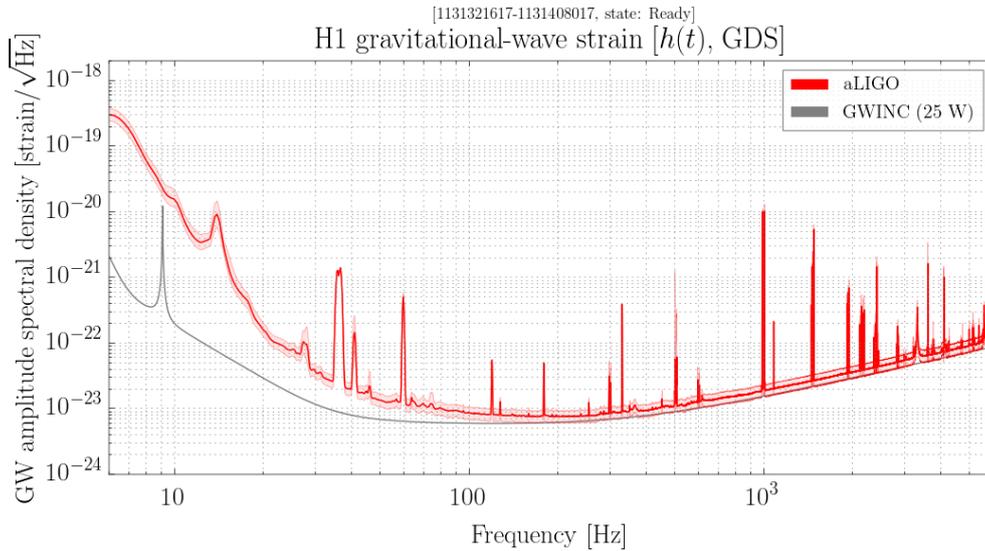


Figure 3.2: Representative noise amplitude spectrum during the O1 period for H1 (red), with aLIGO ideal sensitivity for O1 configuration and 25W laser power superimposed (black). The solid traces represent the median sensitivity and the shaded regions indicate the 5th and 95th percentile over the analysis period. The narrowband features in the spectra are due to known mechanical resonances, mains power harmonics, and injected signals used for calibration [12]. Non-stationarity in the detector noise manifests as variations in the level and shape of these sensitivity curves over time. The grey curve represents the design noise budget for aLIGO during O1. Snapshot taken from November 12, 2015. Generated with [76]

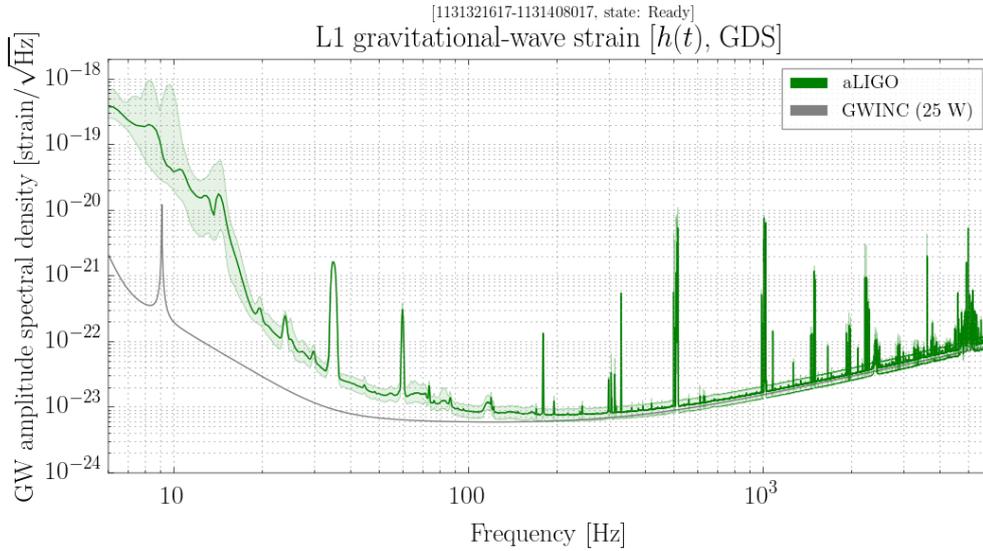


Figure 3.3: Representative noise amplitude spectrum during the O1 period for L1 (green), with aLIGO ideal sensitivity for O1 configuration and 25W laser power superimposed (black). The solid traces represent the median sensitivity and the shaded regions indicate the 5th and 95th percentile over the analysis period. The narrowband features in the spectra are due to known mechanical resonances, mains power harmonics, and injected signals used for calibration [12]. Non-stationarity in the detector noise manifests as variations in the level and shape of these sensitivity curves over time. The grey curve represents the design noise budget for aLIGO during O1. Snapshot taken from November 12, 2015. Generated with [76]

averaged binary neutron star inspiral detection range is the standard astrophysical figure of merit [54]. It is of prime importance in templated-based searches as it characterises the distance to which a typical well-modelled GW signal coming from a binary neutron-star merger can be detected with confidence [55] for a single detector. Figure 3.4 shows this range for both detectors over O1. The two detectors average a distance of about 75 Mpc. H1 was consistently more sensitive than L1, with an angle averaged binary neutron star inspiral range around $\sim 70\text{-}80$ Mpc compared to $\sim 60\text{-}75$ Mpc at L1, indicating comparatively more noise at L1 in the ~ 20 Hz - 300 Hz band.

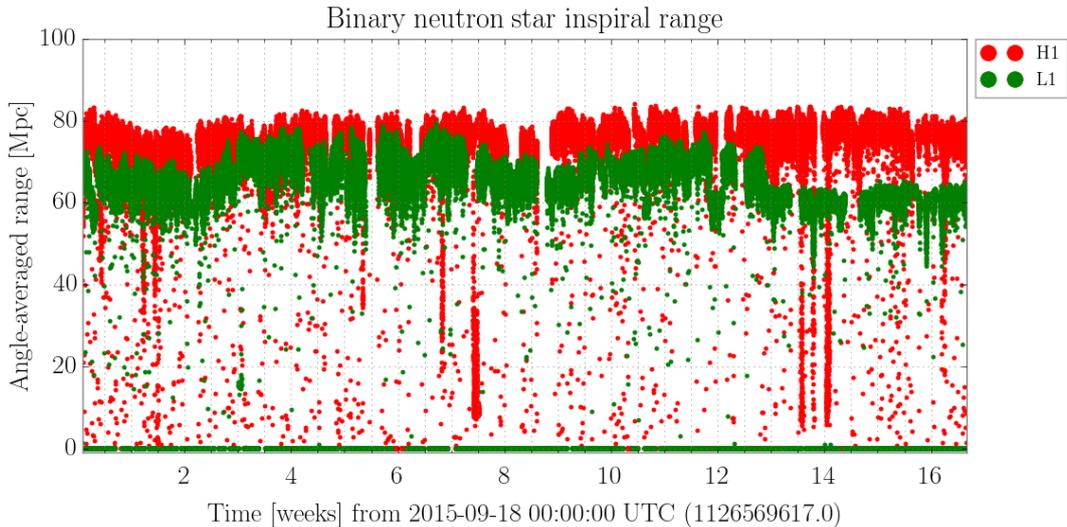


Figure 3.4: Average detection range for binary neutron star mergers using template-based filtering during O1, at H1 and L1. H1 is consistently more sensitive than L1, with a detectable range around 10 Mpc higher.

In the 24 - 1000 Hz band targeted by the search, the calibration uncertainty is less than 19% in amplitude and 10 degrees in phase [25].

3.2.3 Glitches

The detectors are affected by transient non-stationary noise events known as “glitches” that can mimic genuine signals, but have distinct instrumental (or local) environmental origins. Due to the high likelihood of random noise fluctuations manifesting simultaneously yet independently in each detector, many glitches are observed in coincidence and can have properties similar to real gravitational waves like large bandwidth, long duration or time-dependant frequency evolution. During the O1 period, nearly 200,000 glitches were recorded by X-SphRad above SNR of ~ 3 -4.

Glitches with common characteristics are classified in families [107]. X-SphRad is naturally robust against millisecond-long families due to the chosen time-frequency resolution (Section 2.4.3), but is particularly sensitive to second-long glitches regardless of the coherence as the statistic (Eq. 2.10) includes auto-correlation terms. Even though incoherent glitches are eliminated in post-processing by the application of the coherent consistency tests (Eq. 2.23 and 2.24), they have a serious negative impact on the pipeline.

A good example is the wandering line glitch (Fig. 3.5), a long-lived narrow-band glitch wandering around ~ 630 Hz in H1. It is present in the data for the first half of O1, until it was associated with the beat frequency between two oscillators [83]. Preliminary tests on the O1 data showed that while the coherent cuts removed most of these glitches, they were so numerous that the small fraction that survived the cuts were the loudest surviving background events and were reducing the pipeline sensitivity drastically. We applied a frequency-band notching from 590 Hz to 640 Hz over the whole O1 period to remove most of the wandering line glitches from the analysis. Moreover, simultaneous noise perturbations manifested by the other detector can make some fragments of the wandering line appear as coherent power and effectively make the noise-generated trigger peak frequency wander around the line. We note that a better approach to wandering line-like glitches is under testing to avoid taking such a drastic measure in future analyses.

Because X-SphRad only generates a fixed number of triggers per analysis block, loud glitches can *saturate* the clustering and effectively raise the threshold for clustering selection on time-frequency maps (Figure 3.6). When in the vicinity of such glitches, any signal of realistic amplitude present in the block time will remain unclustered and be effectively removed from the analysis. Depending on the

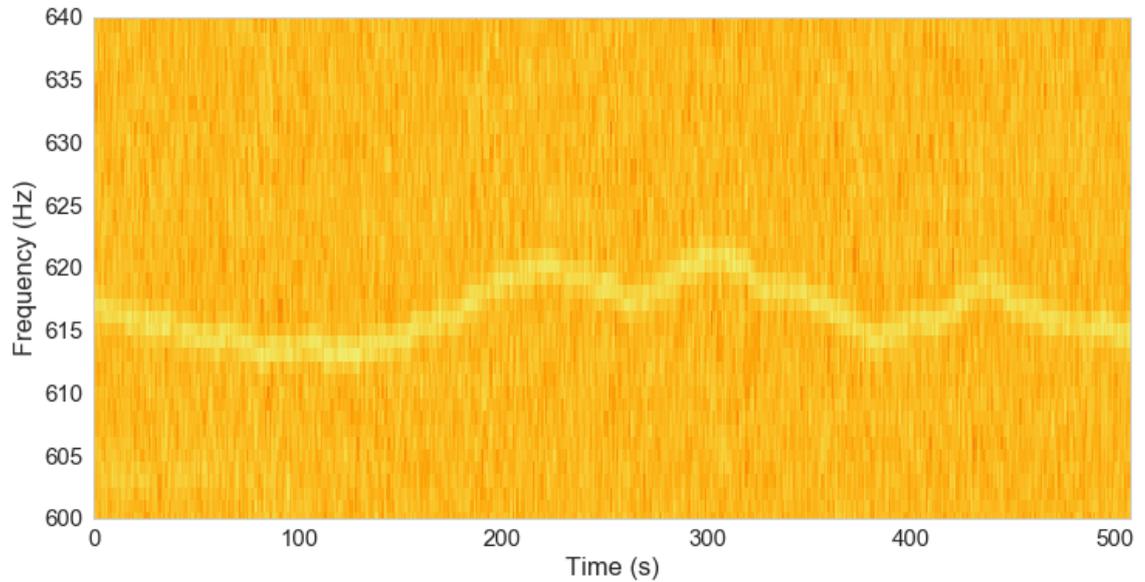


Figure 3.5: Time-frequency map showing the wandering line glitch at ~ 620 Hz. The associated pixels are the loudest connected pixels in the block time. Noisy pixels of comparable amplitude around the wandering line have a tendency to get clustered together and corrupt neighbouring frequencies. A frequency notch between 590 Hz and 640 Hz was applied to remove the effect of this glitch family.

power and frequency variations over time, a single glitch can be split into multiple triggers, each preventing other areas of interest to be followed-up. We suggest a new clustering method in Chapter 4 to increase robustness against this effect.

Another negative impact of glitches is that they can affect the reconstruction of event property if appearing in the same time-frequency space as candidate events. In such cases, the glitch can be clustered with the candidate event and corrupt the reconstruction of event properties such as energy, amplitude or time-frequency properties (Figure 3.7).

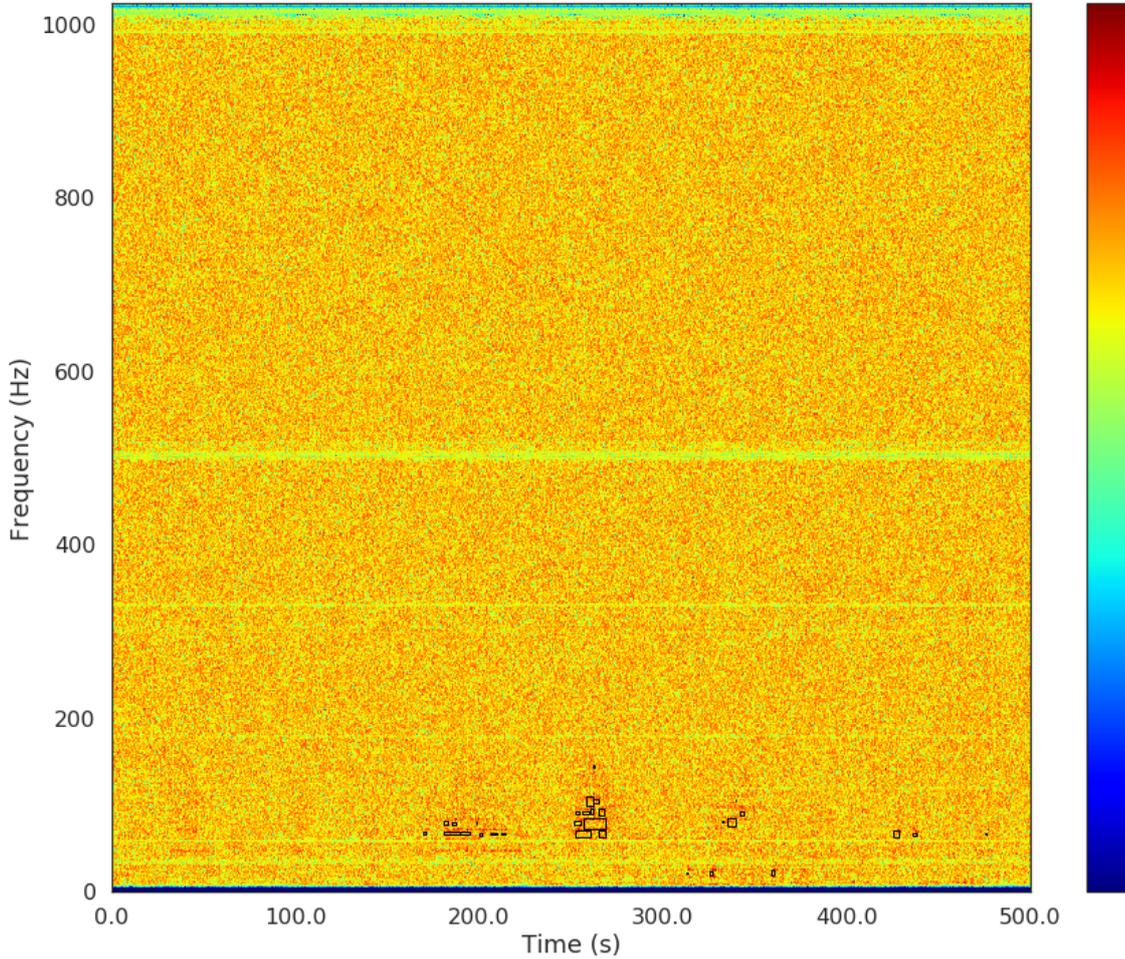


Figure 3.6: Time-frequency map illustrating the effect of glitches on cluster selection. The time-frequency map contains 3 particularly loud glitches at 200 s, 260 s and 360 s as well as a simulated signal from 150 s to 250 s in the 700-900 Hz band. The red lines indicate the bounding boxes of the triggers followed up by X-SphRad. All 30 of the selected triggers are part of the glitches, despite the simulated signal being loud enough to be visible by eye in the data. The white band at ~ 500 Hz is caused by the imperfect whitening of the instrumental violin modes. White grid lines were added to the figure.

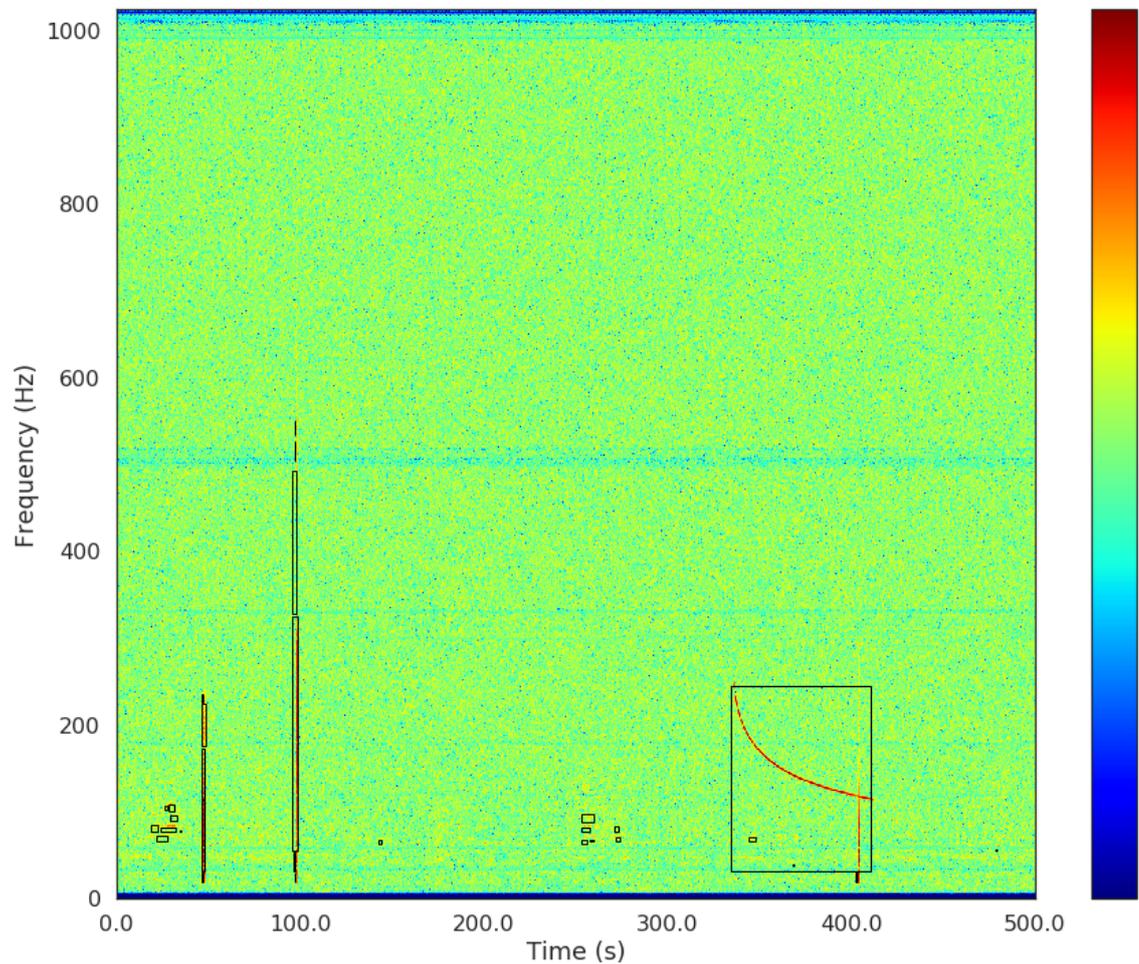


Figure 3.7: Time-frequency map showing the effect of glitches on candidate event reconstruction. Two glitches at 20 s and 260 s, as well as multiple brief broadband glitches at 50 s, 100 s and 400 s (below 100 Hz) are present in the data. There is also a downward-chirping simulated signal from 350 s to 410 s. The red lines indicate the bounding boxes of the triggers followed up by X-SphRad. The short glitch at 400 s gets clustered with the simulated signal, resulting in an overestimation of the signal bandwidth and corruption of the estimated energies.

3.2.4 Data Quality

Periods of poor data quality are identified by examining data from environmental and instrumental auxiliary monitoring channels. We remove these from the analysis times corresponding to such known data quality issues, as defined in [23]. The vetoes applied in this context are divided into categories, depending on the severity of the data quality problem.

CAT1 vetoes are identified when the detectors are not producing data that can be trusted; high noise in the detectors can contaminate an analysis by biasing the power spectrum density estimate of the whitening filter. This data is not analysed and totals to 2.0 days in H1 and 1.4 days in L1.

CAT2 vetoes flag transient excesses of noise in auxiliary channels where there is an understood physical correlation between the auxiliary channels and the detector output. Those short periods are identified using H_{veto} [67]. CAT2 vetoes account for 2.2 days in H1 and 3.9 days in L1.

CAT4 vetoes register times when simulated GW signals (‘hardware injections’) are added to the control feedback loop of one or more optical components to simulate the effect of a GW signal on the detector. They account for minimal time loss, with only 0.1 days in each detector individually.

For completeness, we note that there also exist CAT3 vetoes which identify times where there is a known statistical correlation with the GW channel but the physical connection to auxiliary channels is not understood. These are commonly used when setting upper limits from short-duration GW burst searches [23], but we do not use CAT3 vetoes in this search.

X-SphRad requires 512s of continuous coincident science mode for processing. To

avoid removing an entire 512s block for a veto of only few seconds of total “deadtime”, data corresponding to CAT2/4 times are analysed as normal, but any trigger event found in the GW channel at vetoed times is automatically discarded from the analysis in the post-processing step.

3.2.5 Additional Vetoes

Historically, the data quality vetoes are based on comparing noise transients in auxiliary channels to events reported by short-duration GW burst searches [21] and are thus optimised for searches targeting short duration signals. However, we also consider additional vetoes generated specifically for long duration searches, based on a few channels that were found to be reliable indicators for suboptimal data acquisition in regard to long duration searches. Specifically, a short-duration wide-band glitch of high power was found to be linked to the RF45 modulation [32]² and have been adopted by all the long-duration burst searches in O1. These account for a total dead time of 0.47% of the O1 run in H1 and 1.0% in L1. Table 3.2 shows the list of veto channels.

3.3 Search Parameters

We searched the most sensitive aLIGO frequency band, from 24 Hz up to 1000 Hz. We divided each analysis segment into data blocks of 512s to maximise the recovery of longer signals and avoid edge effects: we found that increasing the block time,

²The burst CAT2 short duration vetoes are also effective against RF45 glitches, but have a much larger deadtime ($\sim 2\%$) that was negatively affecting the O1 long-duration Zebragard search.

Channel name
H1:ASC-REFL_A_RF45_I_PIT_OUT_DQ_veto
H1:IMC-IM4_TRANS_PIT_OUT_DQ_veto
H1:ODC-IMC_WFS_DOF3_YAW_HIGH
H1:LSC-POP_A_RF9_Q_ERR_DQ_veto
H1:SUS-SRM_M1_DAMP_V_IN1_DQ_veto

Table 3.2: Channel names and corresponding deadtime for additional vetoes applied specifically in the long-duration O1 searches. The additional vetoes were suggested by the STAMP-AS team [46]. The burst CAT2 short duration vetoes are also effective against RF45 glitches, but have a much larger deadtime ($\sim 2\%$) that was negatively affecting the O1 long-duration ZebraGard search. Source: private email conversation.

from 256s in previous X-pipeline analyses ([3] for example), reduced the tendency for our target signals to be split between blocks.

We operated in the time-frequency (TF) domain by constructing TF pixel maps (TFmaps) using a $1\text{ s} \times 1\text{ Hz}$ Fourier transform with 50% overlap. The resolution, equal in time and frequency, was empirically found to be good for recovering long-duration signals. As best-guess long-duration models tend to exhibit both monochromatic-like and chirp-like features, an unbiased resolution equal in time and frequency is appropriate for maximum feature visibility. We used a maximum spherical harmonic order $l_{max} = 60$, which was found to be a good trade-off between spherical domain resolution and computational time [50]; see Figure 3.8.

3.4 Clustering

Candidate events are identified as groups or clusters of loud pixels in the time-frequency map. The seed and branch thresholds as defined in Section 2.5 were chosen to be identical to the 5% loudest pixels. The pixels on the time-frequency map with the 95% lowest values were discarded.

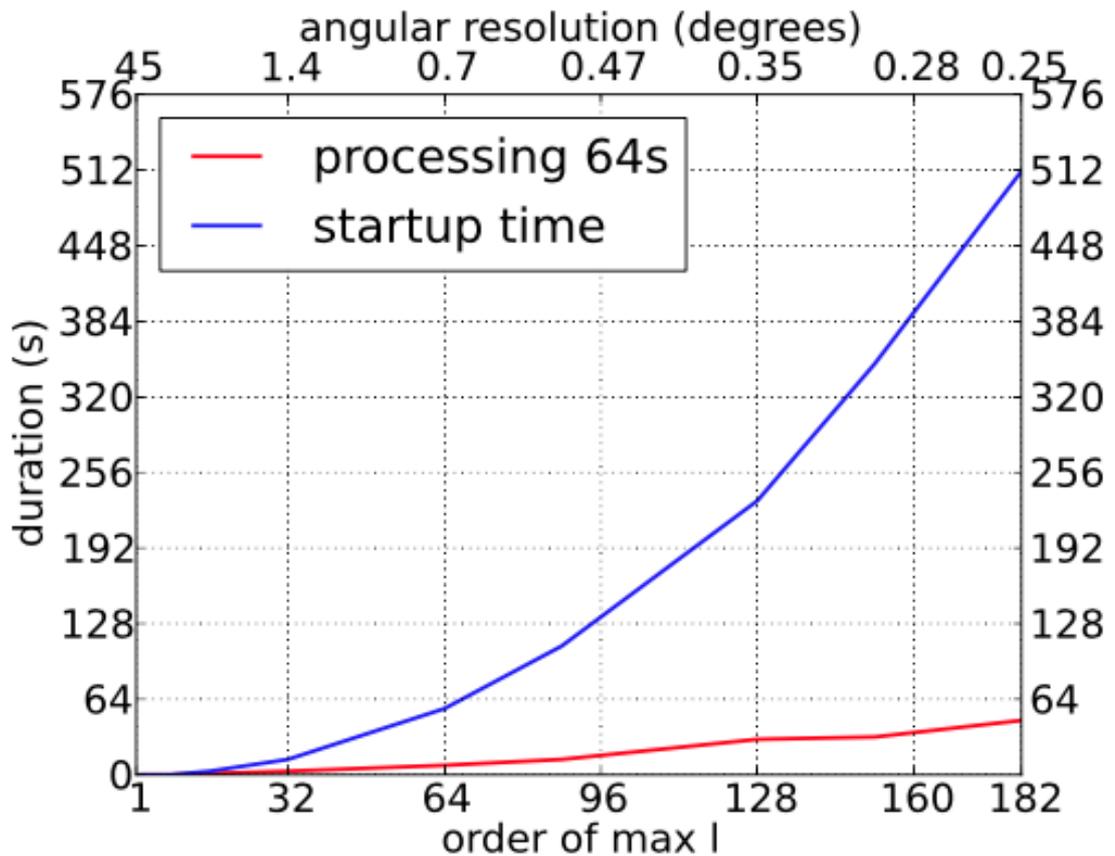


Figure 3.8: Time used to generate the delay matrix and to process 64 seconds of data, as a function of spherical harmonic number l_{max} . X-SphRad has a fast processing time but a significantly longer startup time. The dominant cost is the background analysis using internal timeslides. Each 512 s background job processes 165 circular timeslides (see Section 3.5 for details). With the chosen parameters $l_{max} = 60$, the estimated startup time is approximately 1 minute, compared to about 200 minutes for processing the internal timeslides. The startup time does not significantly affect the total computational cost. Taken from [50]

Long-duration model signals added to the data showed a tendency to be fractured into multiple triggers (Figure 3.9) by various features in the time-frequency map, such as power lines, suspension violin modes, or by the whitening filter overwhitening a frequency band following a loud glitch. In such cases, the clustering procedure returns the injected GW event as several separate triggers, each containing only a fraction of the SNR of the whole signal. Moreover, as only a small subset of pixels are available for energy estimations, the resulting fragments are more susceptible to noise perturbations and being removed during the post-processing by the cuts (Section 2.5.3). The recovery of the full waveform is therefore dependent on a suitable clustering mechanism of the resulting waveform fragments. We implemented an additional *super* clustering step to combine these triggers and reconstruct the whole signal.

The super clustering takes the list of triggers reported by X-SphRad, and selects those whose bounding boxes overlap to within 5 sec and 8 Hz. They are then combined into a single trigger with a resulting bounding box defined by the smallest time-frequency rectangle containing the bounding boxes of all the constituent triggers. A chain of several triggers may be combined into a single trigger in this way. Moreover, the likelihoods assigned to the combined trigger are weighted sums of the likelihoods of the constituent triggers, with weights chosen empirically to penalise broadband long-duration glitches. The weighting factor favours triggers which occupy only a small fraction of their time-frequency box (i.e. a curve) and discriminates against glitch-like triggers as they tend to have a large number of loud

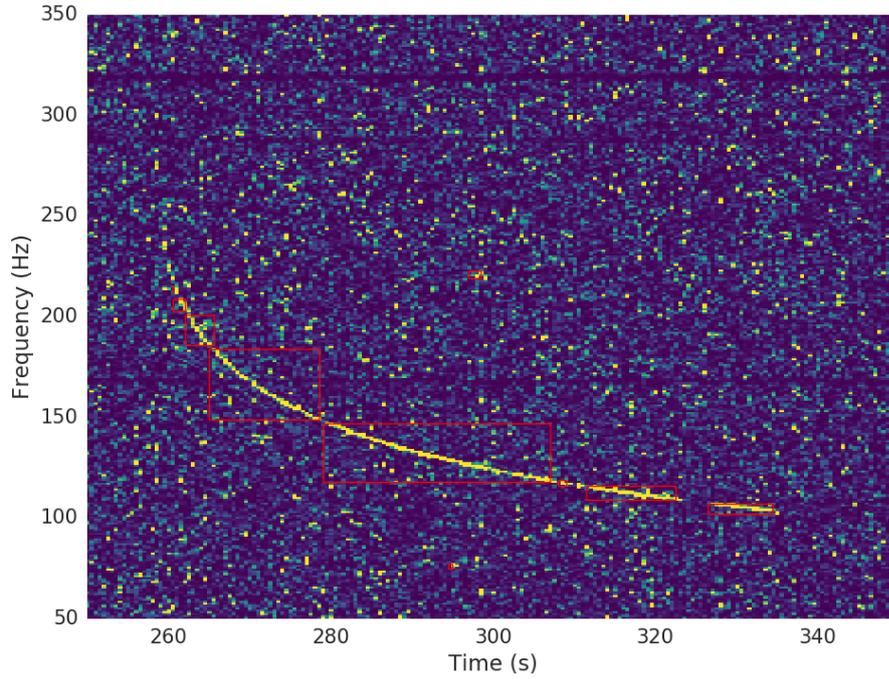


Figure 3.9: Example of long duration signal being fractured into multiple sub-triggers. The signal is split into 8 distinct clusters, each containing only a fraction of the signal total SNR, and will therefore have sub-optimal energy and parameter estimations. The simulated signal is also segmented around 120 Hz due to the effect of the whitening filter.

pixels constrained in a relatively smaller time-frequency volume.³ The weight

$$(3.1) \quad \alpha_{(2\text{det})} = \frac{\Delta T \times \Delta f}{\log(\mathcal{E}_{sl}) \times P^2}$$

is applied to all triggers' background, foreground and simulation trigger likelihoods *before* clustering; here ΔT is the trigger's duration, Δf is its bandwidth, \mathcal{E}_{sl} is the standard likelihood as defined in Eq. 2.10 and P its number of constituting pixels. The weighting factor overall increases the pipeline sensitivity by 40%-120% for astrophysical waveforms (Figure 3.10).

³See Chapter 4 for a discussion of improvements in smart clustering and glitch rejection

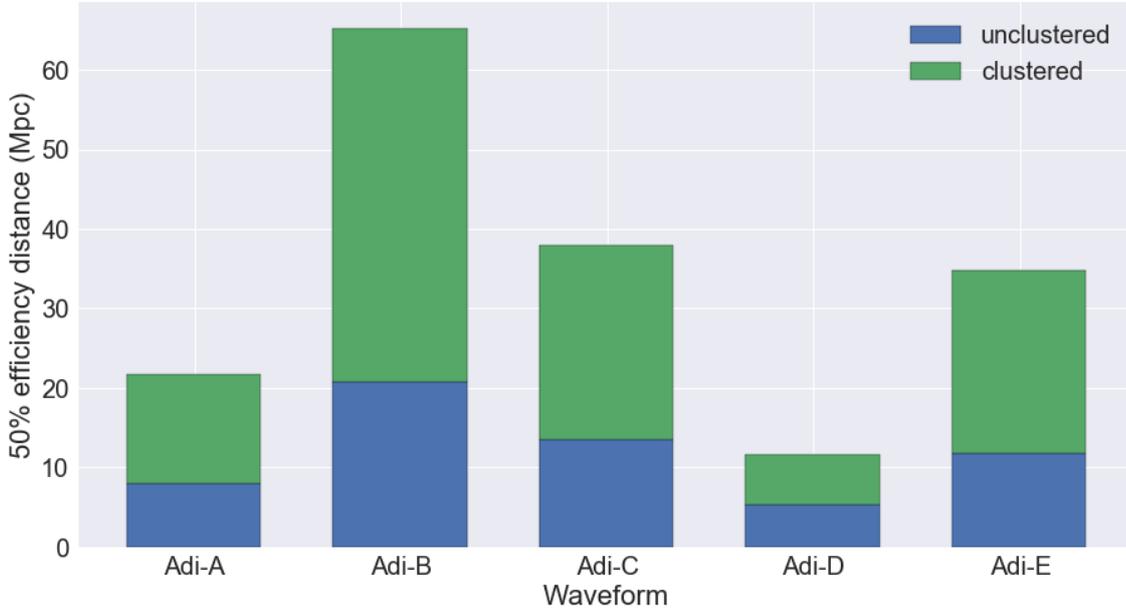


Figure 3.10: Effect of the smart clustering on the 50 % distance efficiency recovery for the astrophysical ADI waveforms. The smart clustering increases the recovery distance by a factor of 2 for Adi-D to 3 for Adi-B.

During the clustering step, we also apply the frequency notching discussed earlier by discarding any trigger with peak frequency between 590 Hz and 640 Hz to circumvent the effect of the wandering line glitches. Moreover, any trigger that is made up of less than 5 pixels is discarded to avoid super-clustering fragmented wide-band long-duration blip-like glitches.

3.5 Background Estimation Procedure

To measure the significance of candidate GW events and differentiate between noise and signal, we estimate the background using the data itself. By using the implementation of the timeslide technique generally adopted in searches for GW transients [9], the data stream from one arbitrarily chosen detector in a network

is shifted by a non-physical time offset and analysed as normal. To prevent the resulting data to be related to a physical source location, the time offset must be at least longer than the light travel time between the detectors. In this search, we applied an offset (or *external lags*) of $\pm 1, 2, 3 \times 496$ seconds to the L1 data-stream, corresponding to multiples of the 512 s block with 8 seconds discarded at each end to account for the whitening filter corruption. X-SphRad also relies on time offsets applied to the data stream within an invocation (or *circular shifts*) in each block, and re-use the computationally expensive delay matrix. The circular shifts were done in 3 s steps, giving a total of 165 different shift in each 512 s block. It is important to note that we are implicitly making the assumption that no GW candidate is present in the dataset. In such cases, the events have been removed from the background trials to allow for astrophysical interpretation [13].

The total number of 825 processed timeslides is smaller than expected from $6 \times 165 = 990$ internal and external slides because the coincident livetime in time lags usually isn't exactly equal to the coincident livetime at zero-lag. As a consequence, the number of data blocks available is typically not exactly a whole number multiple of the number of on-source data blocks. The pipeline set-up procedure rounds down the number of background blocks to be exactly a whole-number multiple of the on-source data blocks. For O1 the number of blocks produced with the 6 *external* lags was slightly less than 6 times that of the on-source blocks, so the number of *external*-lagged blocks was rounded down to 5 times the number of on-source blocks.

We also remind the reader of Section 2.5.1 which states that the background trials are randomly split into two equal groups: one is used to tune the background rejection cuts, the other is used for estimating the False Alarm Rate (FAR) or False Alarm Probability (FAP) of on-source events.

The pipeline automatically tunes the coherent background rejection thresholds on the group used to construct optimal cuts, and maximise the recovery of simulated signals while minimising the amount of significant background events. The automatic procedure selected a value of 2.5 as the optimal veto threshold for the alpha cut, and 1.45 as the optimal value for the fixed coherent ratio threshold. A more detailed explanation of the automatic optimal tuning cut procedure as used in X-SphRad can be found in [95].

Figure 3.11 shows the distribution of trigger significance for both foreground and background, before and after applying coherent cuts. The clear correlation between the off-source and on-source distributions *before* cuts can be attributed to the fact that X-SphRad clustering algorithm triggers on both the auto- and the cross-correlation terms; loud single-detector glitches in the data appear as triggers in all the 165 circular slides. Only the 6 external lags may differ in what loud glitches they report, as those loud glitches fall into or outside of time-shifted coincidence segments. The fluctuations between the on-source and off-source distributions should be interpreted in terms of having at best a few independent trials. This is particularly evident for the very loudest glitches around significance 10^8 , where the on- and off-source distributions change in lock-step due to individual loud glitches. As an example, the blue curve (pre-vetoing background) ends at a rate that corresponds to 4 events; i.e., the loudest pre-vetoing event showed up in 4 of the 6 external lags. However, we stress that the pre-veto distribution is not used by X-SphRad and only the distribution after cuts is used to identify potential GW events.

On condition that the time shifts are non-physical and independent, the post-cut triggers generated by this method should behave like independent realisations of the background and allow for statistically valid estimation of noise. We test by

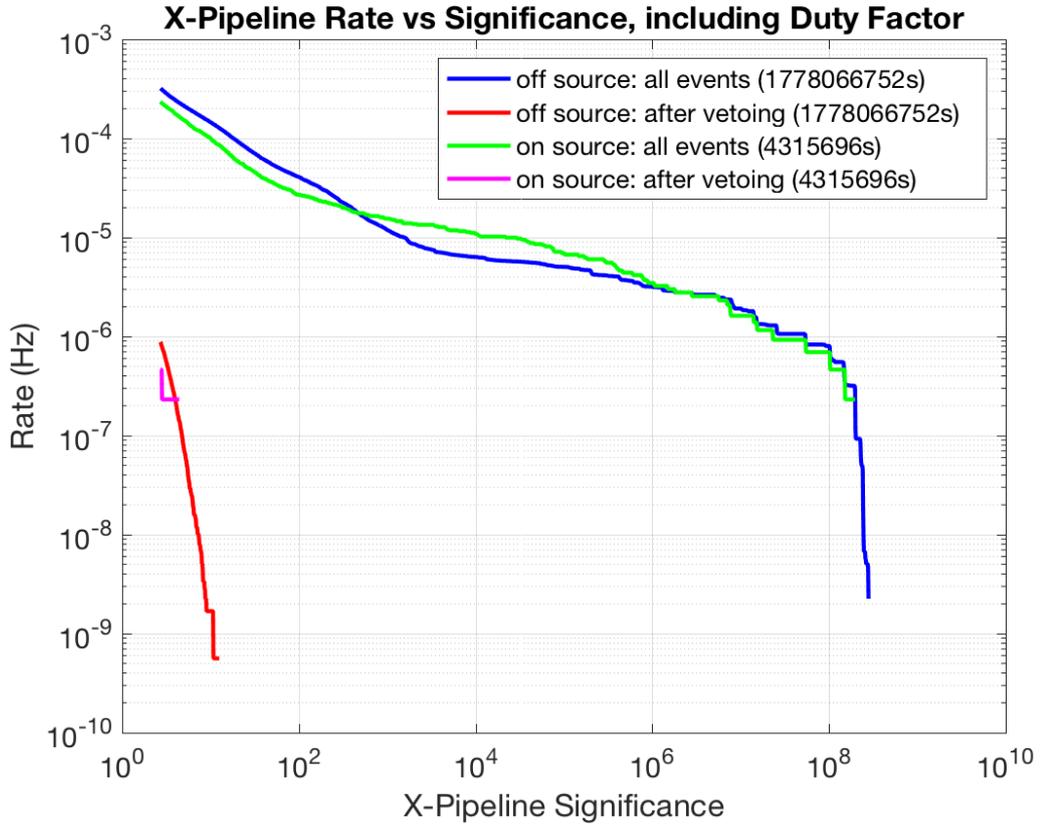


Figure 3.11: Histogram of trigger rate versus significance threshold, including triggers from all upper-limit background jobs and triggers from all on-source jobs, before and after all vetoes & background rejection cuts. There is a clear correlation between the off-source and on-source distributions before cuts, as discussed in the text.

comparing to the Skellam distribution test as applied in the O1 black hole binary search[19]; assuming each time lag is an independent trial, the number of triggers per lag should follow a Poisson distribution with the same mean μ . It can then be shown that the difference in the number of triggers in two consecutive trials, $N_{i+1} - N_i$, should follow a Skellam distribution with mean 0 and variance 2μ . As seen from Figure 3.12, the measured distribution of $N_{i+1} - N_i$ closely follows the Skellam distribution. We conclude that the post-veto distribution is fit to be used to

rank GW candidates.

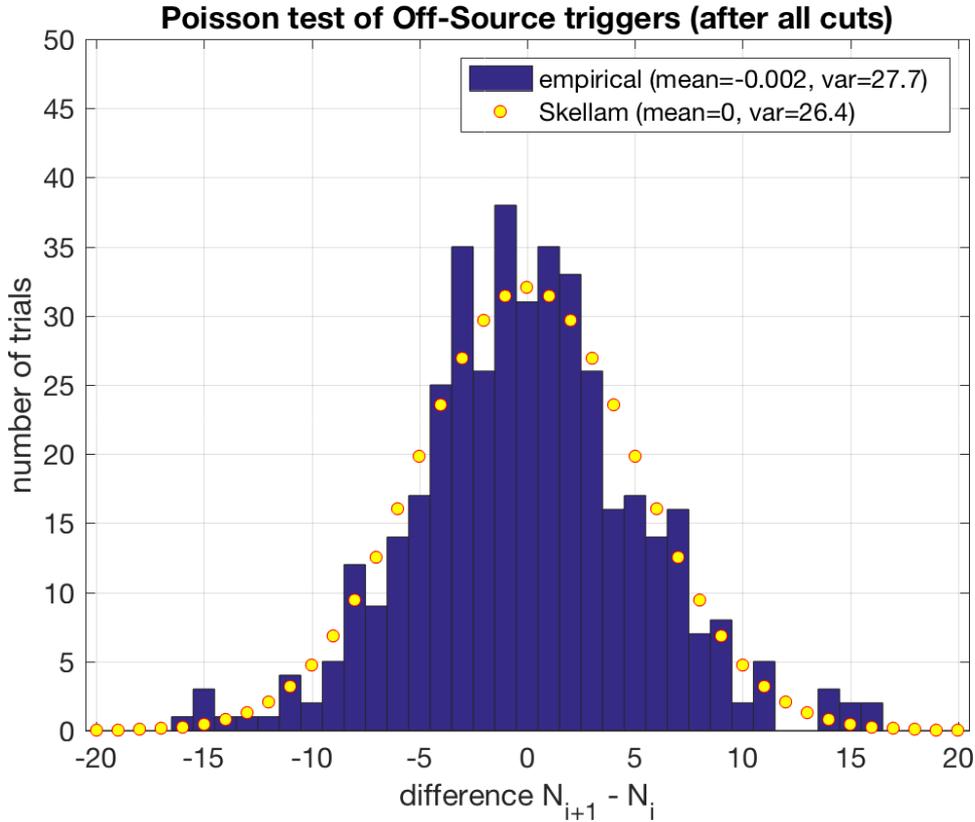


Figure 3.12: Skellam test for the off-source triggers surviving all cuts comparing the empirical distribution of $N_{i+1} - N_i$ to the Skellam prediction. The empirical distribution of the background triggers after cuts in blue is consistent with the Skellam distribution in yellow.

3.5.1 Properties of Background Triggers

Because the selection statistic includes both the cross- and auto- correlation terms (Eq. 2.18), glitches do not require they be coherent to generate a trigger⁴. As a consequence, single-detector glitches can be coherent with relatively quieter glitches

⁴Figure 3.5 is a clear example of the case of a single detector glitch appearing in the coherent search.

- or even with Gaussian background noise fluctuations - in the other detector, and generate a trigger.

Background triggers tend to have very little correlation between detectors. As a consequence, their coherent energies tend to be approximately equal to their incoherent energies. Large fractional excesses of coherent power are only seen in the lower energy events (Figure 3.13).

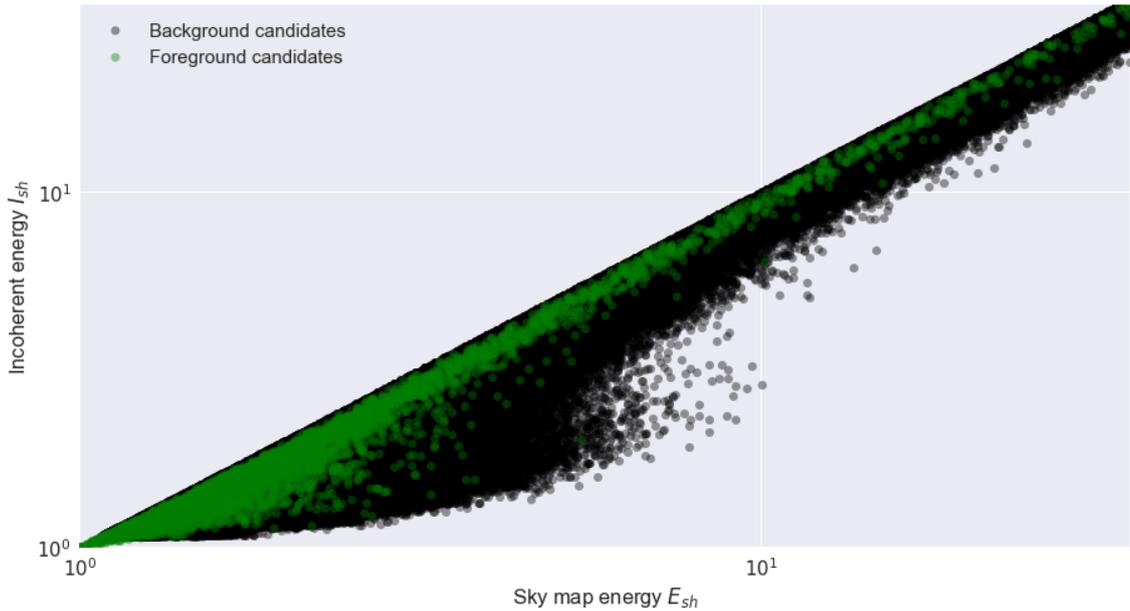


Figure 3.13: Background and foreground candidate events E_{sh}/I_{sh} energy plane. Low energy candidate events (< 10) tend to have larger sky map energy E_{sh} than incoherent I_{sh} energy. This observation motivates the alpha cut (Eq. 2.24).

The distribution of background triggers is not uniform in the frequency space (Figure 3.14); most of the background triggers pre-veto lie between 24 Hz and 300 Hz, with only a marginal fraction having peak frequencies above 600 Hz. This observation indicates that low-frequency noise fluctuations are more frequent than their high-frequency counterparts and are thus more likely be chosen as follow-up potential triggers in the clustering phase.

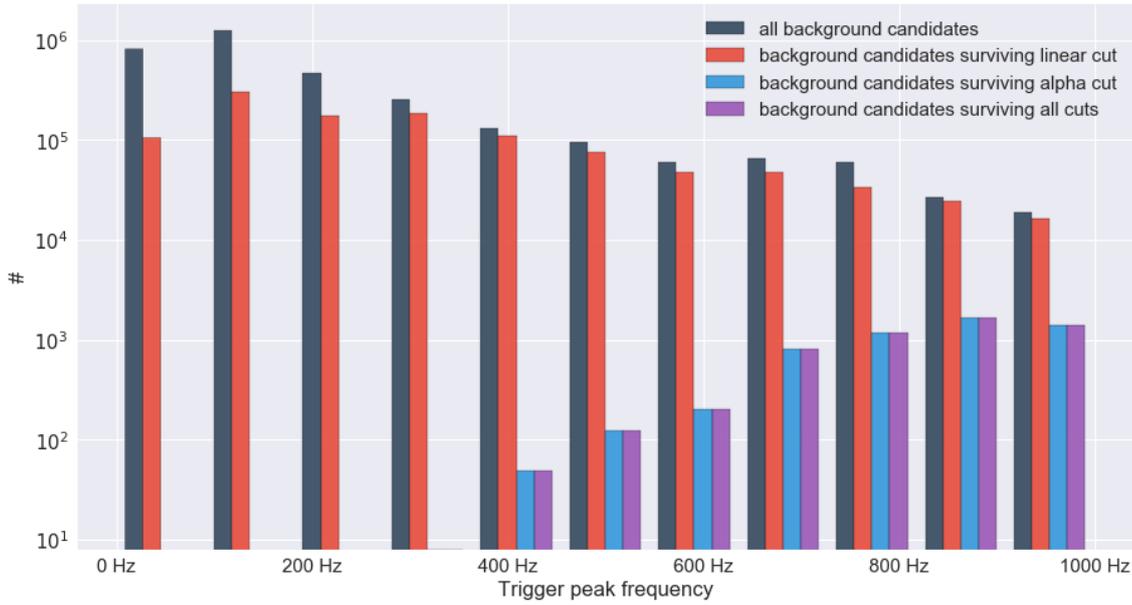


Figure 3.14: Distribution of background triggers before cuts, after the linear and alpha cut as well as final distribution after all cuts vs peak frequency. Most triggers are below 300 Hz before cuts. The linear cut does not seem to discriminate based on the frequency, as the distribution post linear cut follows the original distribution. The alpha cut exhibits a frequency-biased selection effect with most of the surviving triggers located above 600 Hz. The final distribution after all cuts follows the alpha cut distribution, indicating the cut to be the most restrictive overall.

However, the linear cut kills mainly high energy events (> 200), and does not seem to discriminate against frequency, with the pre- and post-cut distributions agreeing reasonably well. Overall, the fixed coherent ratio threshold is the most restrictive and has a bias for high frequency triggers. After applying the cuts, the vast majority of the surviving background triggers have a peak frequency higher than 600 Hz. The alpha cut is aimed at low energy events (< 200) which seems to favour high-frequency events. Further investigation on the energy ratio of background triggers before cuts (Figure 3.15) shows that the vast majority of low frequency triggers tend to have a low alpha ratio.

The correlation between background trigger properties before all cuts (Figure

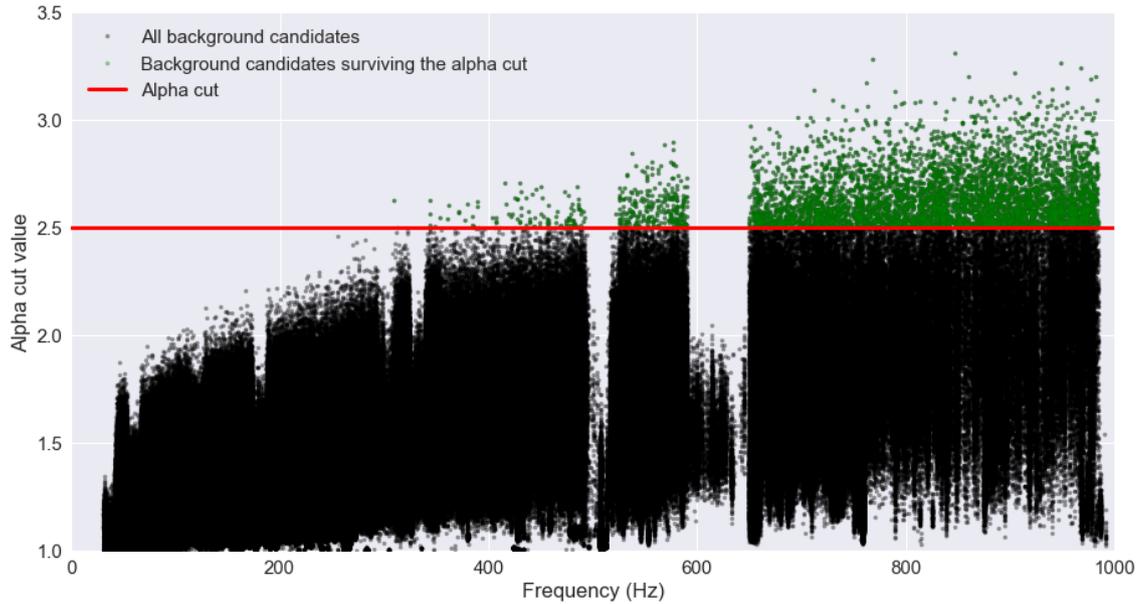


Figure 3.15: Alpha ratio (Eq. 2.24) vs peak frequency for background candidates. The vast majority of low frequency triggers tend to have a low alpha ratio and are thus more likely to be rejected by the alpha cut.

3.16) exposes several relationships. As expected, the number of pixels in a trigger is correlated with its duration and bandwidth, but also has a high correlation with the trigger peak frequency. As the trigger bandwidths and durations increase as a function of peak frequency, this indicates that longer and wider glitches are more frequent at higher frequencies. The high correlation between the number of pixels and the energies is expected, as the sky map energy and incoherent energy are computed from the same number of pixels in the trigger.

After all cuts (Figure 3.17) however, shorter candidate events are more likely to exhibit high sky map energy E_{sh} and incoherent energy I_{sh} . As expected, the bandwidth is still correlated with the number of pixels.

In total 5443 background events passed the cuts and were thus recognised by the pipeline as being consistent with a GW signal. As they were generated using

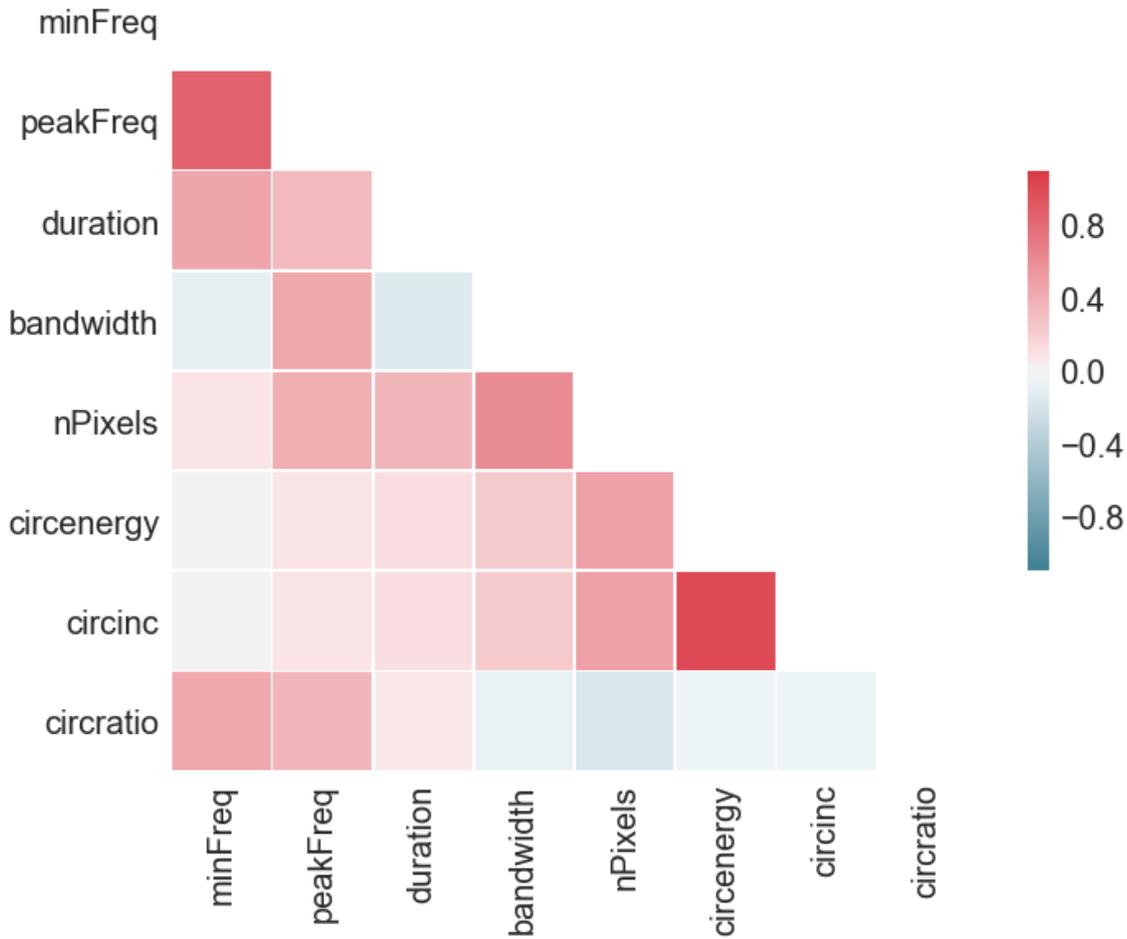


Figure 3.16: Correlation matrix between properties of background triggers before all cuts. The number of pixels is correlated with the duration and bandwidth, themselves correlated with the trigger peak frequency. The effect of the weighting factor during the clustering phase are seen with the relatively high correlation between number of pixels, bandwidth and duration with the sky map energy E_{sh} and incoherent energy I_{sh} .

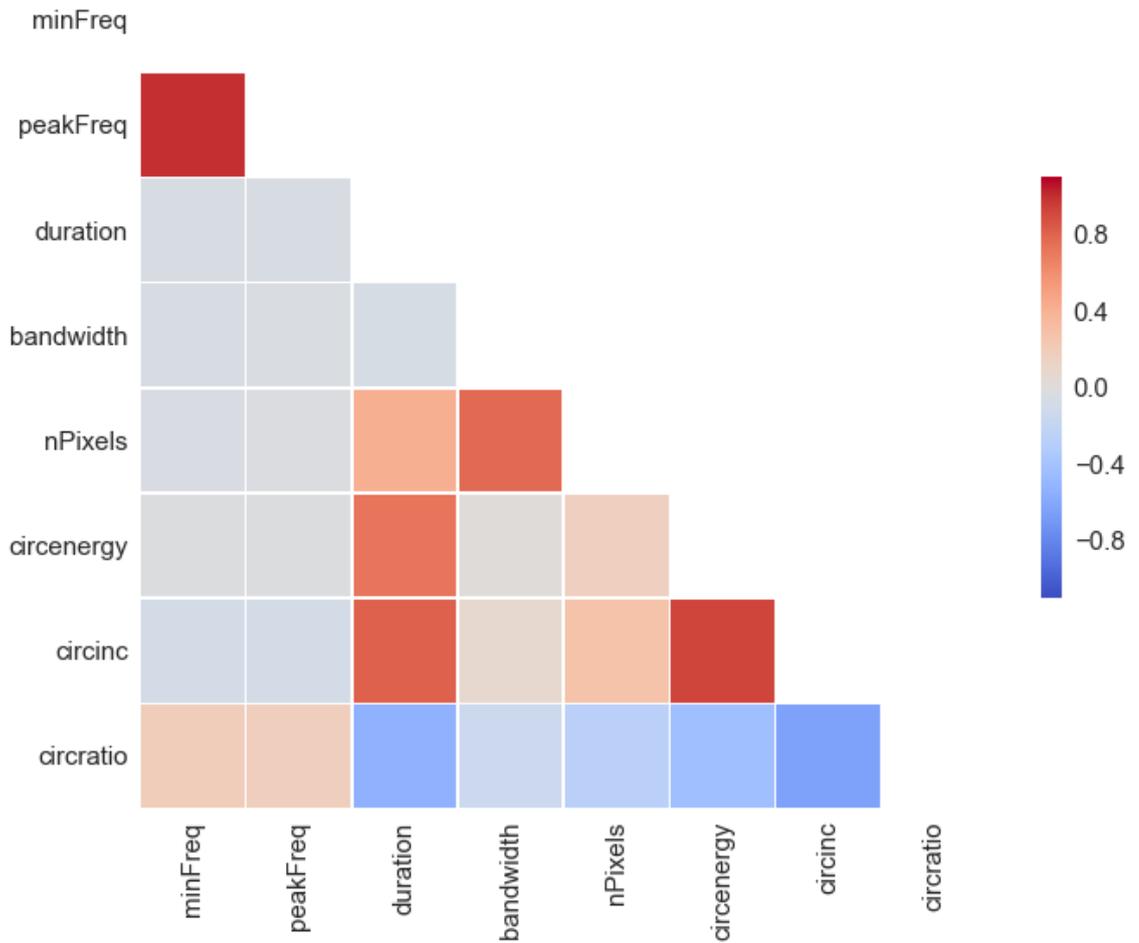


Figure 3.17: Correlation matrix between properties of background triggers after all cuts. The number of pixels is correlated with the duration and bandwidth, themselves correlated with the trigger peak frequency. The effect of the weighting factor during the clustering phase are seen with the relatively high correlation between number of pixels, bandwidth and duration with the circ energy and circ inc.

Table 3.3: Top 5 loudest surviving triggers in background.

GPS for H1	Significance	Peak Freq. (Hz)	Bandwidth	nPixels	I_{sky}	Ratio
1130833954.0	12.66	657.5	14.0	167	5.14	2.46
1132079497.0	10.80	527.5	18.0	27	4.03	2.67
1130812224.5	10.70	650.0	13.0	128	4.13	2.59
1127064560.0	9.02	927.0	15.0	20	1.85	4.86
1126931309.0	8.71	732.0	17.0	21	2.23	3.90

non-physical time shifts, all the triggers can be safely regarded as non-GW signals and representative of the noise background over the course of the O1 period. Table 3.3 shows the 5 most significant background candidate events that passed the fixed and alpha cuts.

Out of the surviving background triggers, the loudest and most significant is shown as seen by X-SphRad in Figure 3.18 with a zoom-in on the trigger in Figure 3.19. The trigger originated from the wandering line at ~ 630 Hz, clustered together with a short glitch at 325 s. The resulting peak frequency moved slightly out of the vetoed frequency region from ~ 590 Hz to ~ 640 Hz. The trigger is clearly not due to a GW.

3.6 On-Source Triggers

The analysis reports 9 potential triggers that could correspond with GW signals. The triggers passed both the coherent threshold and the alpha cut. However, the ranking statistics assigned to the events are relatively low and the triggers are consistent with background noise (Figures 3.20 and 3.11).

Foreground triggers are spread throughout the O1 period, with a slight concentration over the last few months of the run that can be attributed to randomness. Table

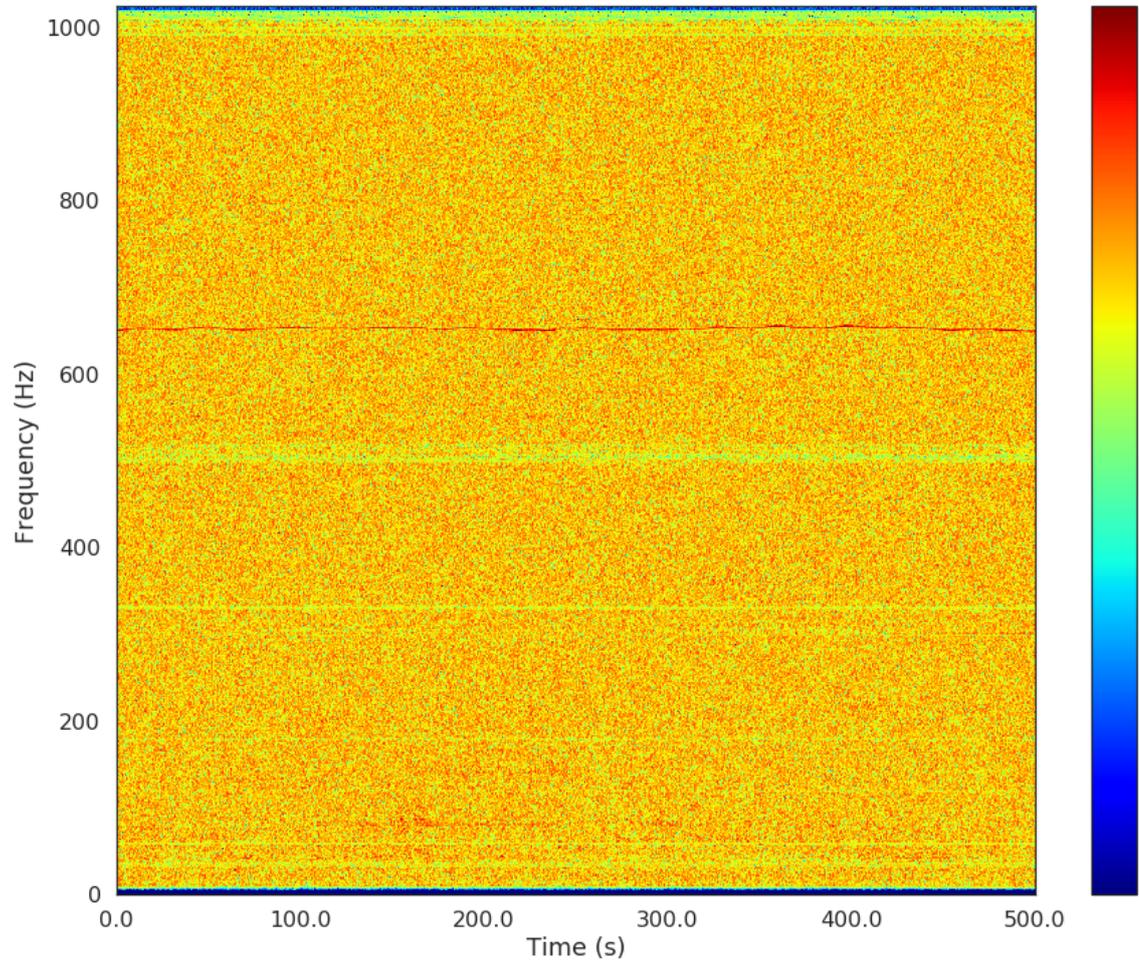


Figure 3.18: Loudest background event surviving cuts. The ~ 630 Hz wandering line at H1 was clustered with a short glitch at 325 s. The resulting cluster peak frequency is slightly above that of the frequency notch applied in the clustering phase. See 3.19 for a zoomed in version.

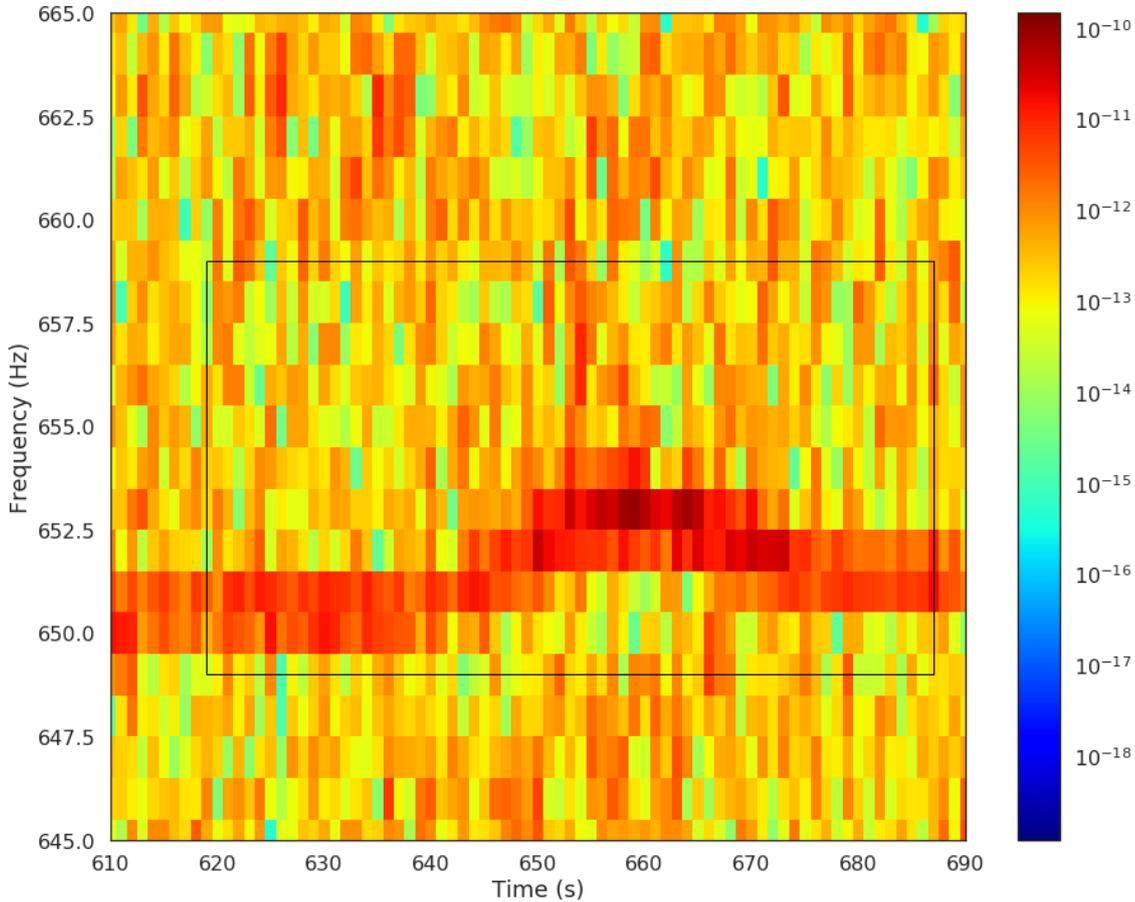


Figure 3.19: Zoomed version of Figure 3.18. A coherent set of pixels from 645 Hz to 655 Hz at 325 s is clustered with the wandering line, increasing the central frequency above the upper cutoff frequency of the notch. The red line is the bounding box of the candidate as seen by X-SphRad.

3.4 shows recovered parameters for the surviving foreground triggers. Excitations in auxiliary channels such as IMC, SUS, and PEM were observed in coincidence with all triggers reported. Figure 3.21 shows the time-frequency maps of calibrated data for the H1 and L1 detectors around the time of the loudest on-source trigger. As there is no clear sign of the event in the single-detector data due to the low SNR spread over second-scale duration, further investigations in the time-frequency domain were not possible. We conclude that we find no evidence for a long-duration

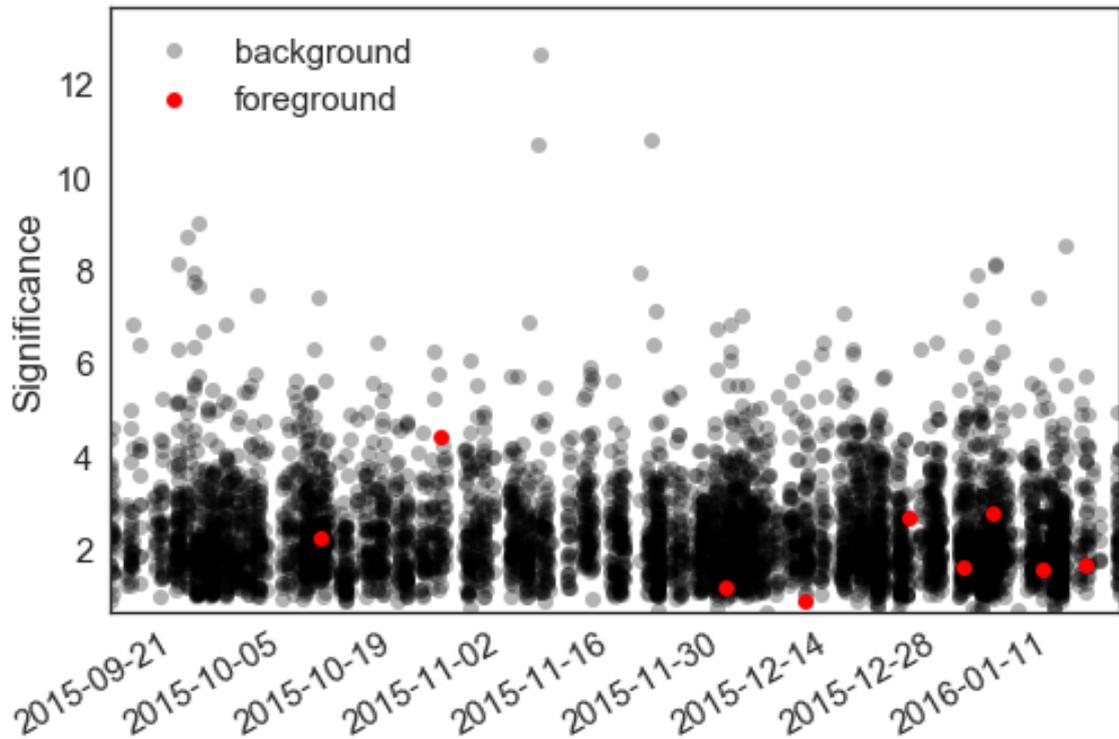
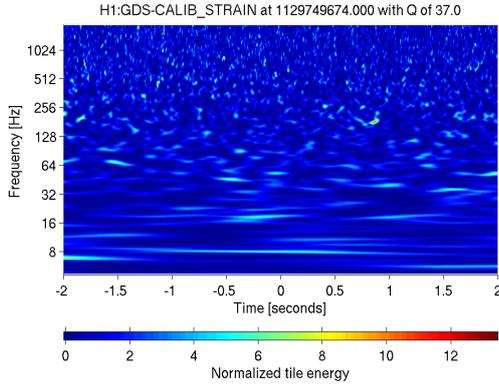


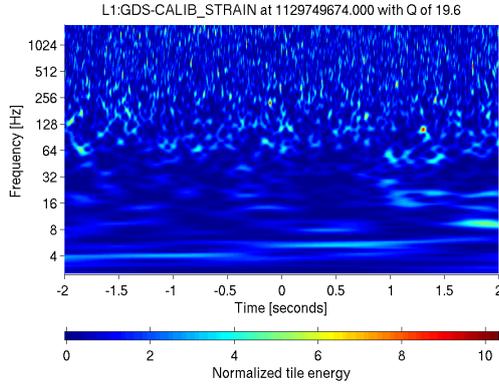
Figure 3.20: Surviving background (black) and foreground (red) triggers during the O1 period, after coherent cuts are applied. The triggers are ranked by significance. Foreground triggers are spread throughout the O1 period, with a slight concentration over the last few months of the run. All foreground triggers are consistent with the background.

gravitational-wave transient during the first advanced LIGO observing run.

Albeit present in the data, the gravitational wave events GW150914 [13] & GW151226 [14] and the candidate event [18] were not recovered by X-SphRad. Due to search time-frequency resolution constraining the duration uncertainty to 1 second, short-duration waveforms can only be recovered if their coherent energies are louder than the noise fluctuations within 1 second at each frequency bin. In such situations, the resulting pixel-cluster is required to have a energy statistic amongst the top 30 clusters to be selected for follow-up. After that, the smart-clustering step



(a) Time-frequency map of the loudest surviving foreground candidate in H1. No particular feature is visible, as expected from the low SNR, and low number of pixels. The trigger is located around 902 Hz in the centre of the spectrogram.



(b) Time-frequency map of the loudest surviving foreground in L1. The trigger is located around 902 Hz in the centre of the spectrogram.

Figure 3.21: Loudest surviving on-source event

Peak GPS time (s)	Significance	FAP	peak Freq (Hz)	Duration (s)	Pixels	Approx SNR
1129749676	4.4	0.43	902	4.0	17	5
1135855931	2.8	0.97	806	2.5	14	4
1134913112	2.7	0.98	842	2.5	13	4
1128421168	2.2	1.00	695	2.5	15	3
1136863671	1.7	1.00	972	2.0	16	2
1135534408	1.6	1.00	753	3.0	14	3
1136392610	1.6	1.00	940	3.0	15	2
1132887849	1.2	1.00	731	1.5	15	2
1133764269	0.9	1.00	663	1.5	19	2

Table 3.4: Surviving foreground candidate events. All were consistent with the background.

discards clusters formed by 5 pixels or less, making the recovery of short duration events of reasonable amplitude by the pipeline unlikely.

3.7 Search Sensitivity

Simulated signals are used to tune the background rejection cuts for optimal sensitivity. However, they are also commonly used to measure the efficiency of burst searches, giving the distance at which a pipeline can reliably detect a simulated signal drawn from an isotropic distribution of sky locations and orientations. The efficiency at a given distance is defined by the fraction of signals at that distance passing all background rejection cuts, and with a ranking statistic equal to or larger than a value corresponding to the given FAR estimated from the background noise.

We injected 16 types of simulated signals covering the target parameter space (Figure 3.22), with a slight concentration in the more sensitive 30 - 300 Hz band. Each waveform was injected ~ 800 times at randomly chosen times over the run period, each with 26 amplitude factors, corresponding to distance, logarithmically spaced from 100 kpc to 1000 Mpc. The source location (ϕ, θ) on the sky and orientation relative to the earth were distributed isotropically as follows :

- $\cos(\theta)$ uniform on $[-1, 1]$
- ϕ uniform on $[0, 2\pi)$
- $\cos(\iota)$ uniform on $[-1, 1]$
- ψ uniform on $[0, \pi)$

The simulated signals can be divided into 2 categories, astrophysically motivated and *ad-hoc*, each further subdivided by families sharing common properties. See Tables 3.5 and 3.7. Additional detection efficiency figures for each waveforms are available in Appendix A.

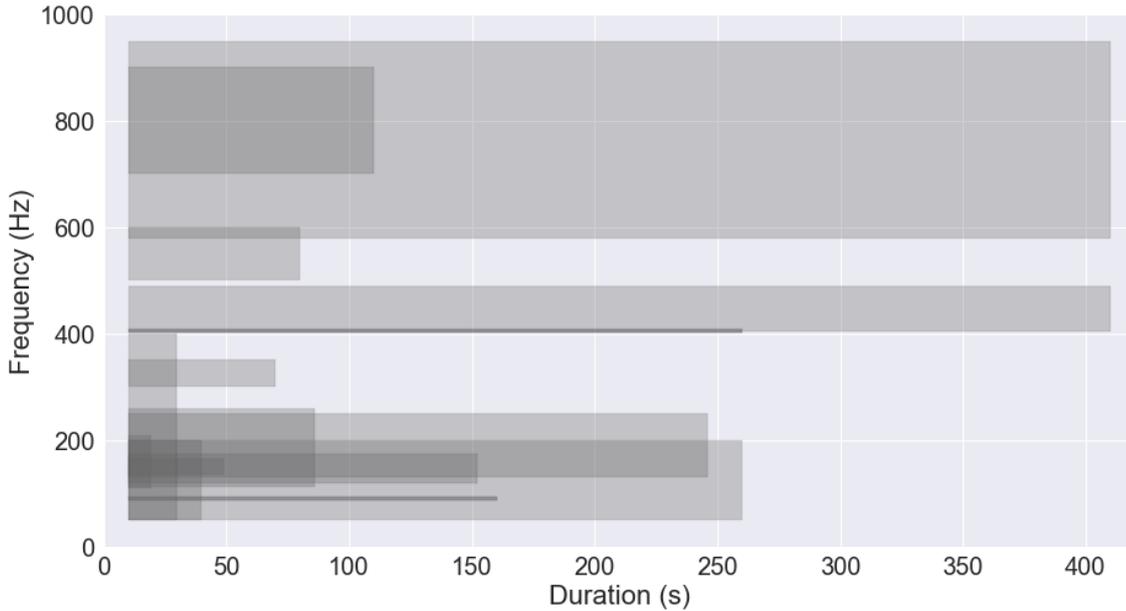


Figure 3.22: Illustration of the parameter space covered by the full set of simulated signals used during the O1 analysis. Each shaded rectangle represents the duration and frequency range of one of the 16 waveform types simulated. The whole frequency space is covered with a slight concentration at lower frequencies in aLIGO’s most sensitive band.

3.7.1 Astrophysical Models

The astrophysically motivated waveform generation mechanisms have been discussed in detail in Chapter 1. They include the Accretion Disk Instabilities (ADI) family and the magnetar family. They are typical, order-of-magnitude approximation waveforms unsuitable for match-filtering searches. Table 3.5 summarises the time-frequency properties of the waveforms used in this search.

Figure 3.23 compares the 50% and 90% efficiency distance for the astrophysical waveforms. Overall, X-SphRad is more sensitive to the waveforms of the ADI family to a distance that is more than an order of magnitude larger than the distance for the magnetar family. Aside from the nominal amplitude definition, differences

Waveform	Duration (s)	Start Freq. (Hz)	End Freq. (Hz)
Adi-A	39	135	166
Adi-B	9	110	209
Adi-C	236	130	251
Adi-D	142	119	173
Adi-E	76	111	260
Magnetar-F	400	579	950
Magnetar-G	400	405	490

Table 3.5: Time-frequency properties for astrophysical waveforms. The ADI family covers the lower end of the targeted frequency space with \sim minute long signals, while the magnetar family covers the middle and higher end with longer duration. Individual TFmaps of each waveform can be found in Appendix B.

between waveforms can be explained by considering their individual energies spread in varying volumes in the time-frequency space.

To estimate the detectability of the waveforms, we can construct an optimal SNR as the SNR needed to recover 50% of the waveforms with face-on orientation in a single aLIGO detector in noise estimated from aLIGO design sensitivity, weighted by a factor of 3 to match current aLIGO sensitivity and assuming a perfect match-filtering search. Table 3.6 shows that the optimal SNR grows with duration and bandwidth, because the waveform energy is spread over a longer period in the time-frequency space where noise accumulates.

Following [29], we can construct a more adapted figure of merit representing the detectability of the waveforms that takes into consideration the time-frequency spread. By creating a time-frequency volume

$$V = \delta t \delta f$$

where δt is the waveform duration and δf is the waveform bandwidth, [29] predict that detectability in Gaussian noise should vary with the volume weighted SNR as

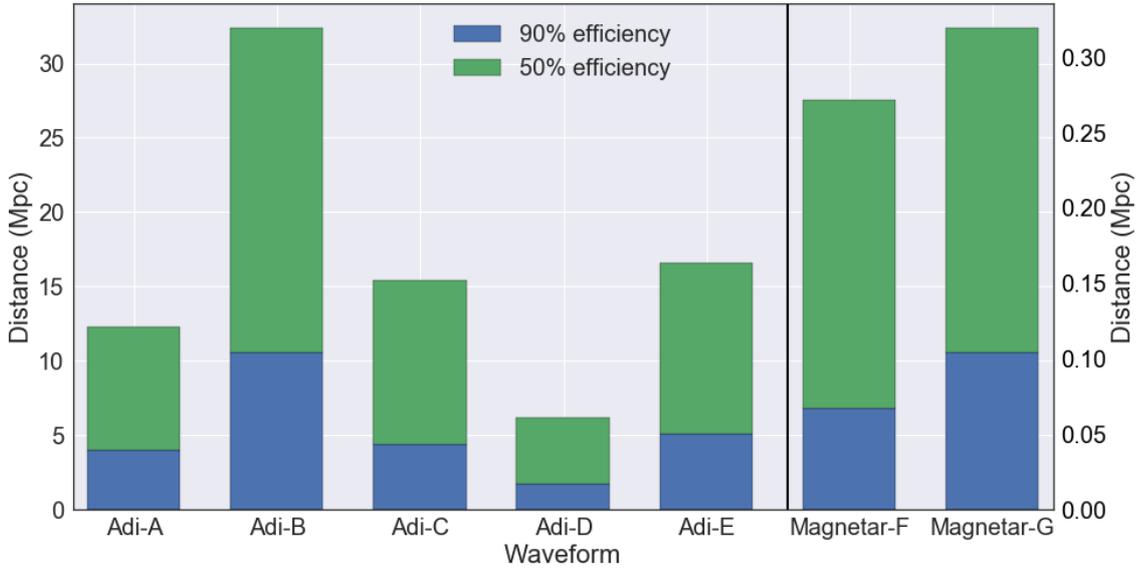


Figure 3.23: Comparison of 50% efficiency and 90% sensitivity distances of X-SphRad during O1 for astrophysical waveforms. Differences between ADI waveforms can be explained by the spread of power in longer durations and wider bandwidths. The differences between the ADI family and the magnetar family are due to the magnetar lower intrinsic amplitudes.

$$(3.2) \quad \text{Detectability} = \frac{SNR}{V^{1/4}}.$$

where V is the volume in the time-frequency plane. We use this measure to rank the inherent detectability of the waveforms while taking into account the respective energy spread. Table 3.6 shows that all astrophysical waveforms are within $\pm 15\%$ of each other, indicating that the performance of X-SphRad is approximately consistent across the astrophysical waveforms tested.

The recovered parameters for the ADI waveforms are consistent with the injected parameters (see for example, Figures 3.25 and 3.26). The bandwidth is overestimated at short distances then underestimated onward. As the ADI waveforms start with high amplitude then chirp down with lower energy (Figure 3.24), the clustering

Table 3.6: Detectability of the astrophysical waveforms. The optimal SNR is the SNR needed to recover 50% of the waveforms in a single aLIGO detector in noise estimated from aLIGO design sensitivity, weighted by a factor of 3 to match current aLIGO sensitivity. For this estimation, the waveforms were generated assuming face-on orientation. Taken from [29]

Waveform	Dist 50% (Mpc)	Optimal SNR	Detectability Statistic 3.2
Adi-A	12.3	314	118
Adi-B	32.4	418	143
Adi-C	15.4	915	132
Adi-D	6.1	636	140
Adi-E	16.6	687	133
Magnetar-F	0.3	1012	140
Magnetar-G	0.3	1641	158

algorithm needs a few loud pixels from waveforms to start up. Once the waveform is loud enough to be detected by the clustering algorithm, the effect of Fourier leakage can be seen and the bandwidth is overestimated. The effect is expected as the higher injection scales are extremely loud, with $\text{SNR} > 1000$ which saturates the whitening filter. The duration recovery shows a similar effect, mitigated by the relative constant spread of energy over time of the family.

The effect of the Fourier leakage on the time-frequency map is visible in Figure 3.27 where the amount of mean pixels recovered grows with closer distances. Such distances are unrealistically close for an astrophysical event, thus the Fourier leakage does not impact the search.

The magnetar (Figure 3.28) waveforms exhibits a slow ramp up of power in both time and frequency. Combined with its quieter definition amplitude, the pipeline does not fully recover the waveform in duration (Figure 3.30) nor in bandwidth (Figure 3.29).

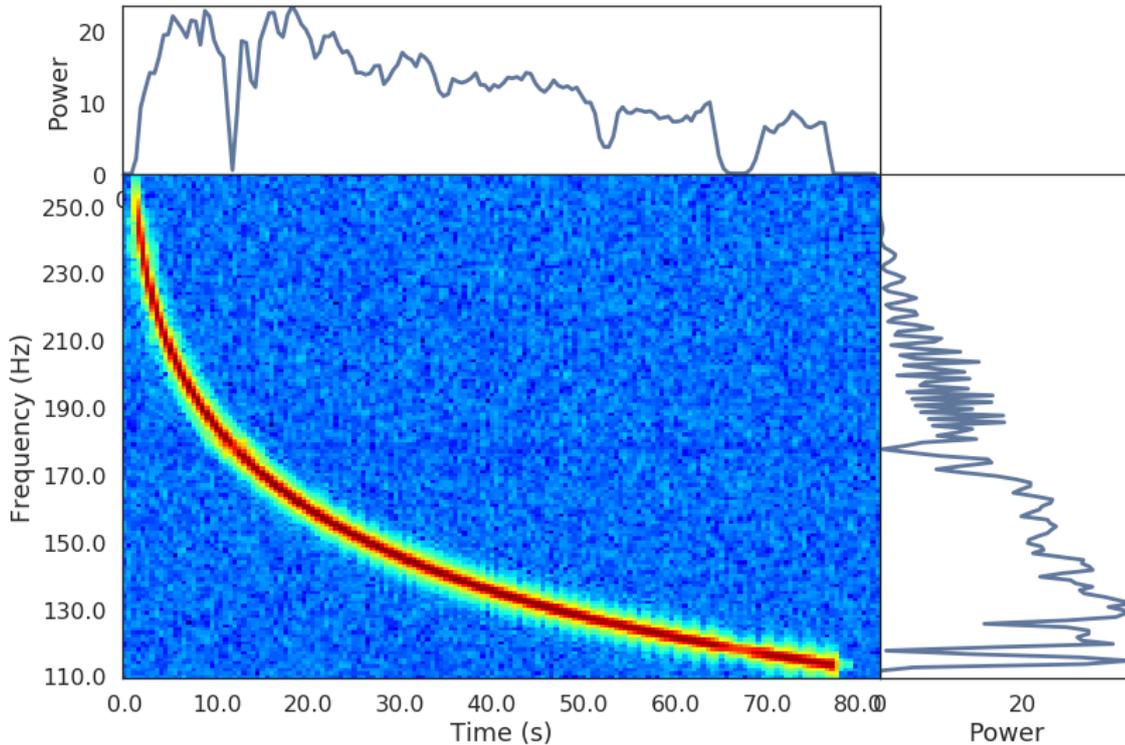


Figure 3.24: Energy spread of ADI-E in time and frequency. The waveform energy over time is slowly decreasing, while the spread of energy in frequency shows an inverse chirp with power concentrated towards the end of the waveform. This is due to the monochromatic-like behaviour of the waveform at the end of its lifetime.

3.7.2 Ad-hoc Models

Ad-hoc models are intended to cover the parameter space with easily understood waveforms and test the overall pipeline search capabilities. They are not generated from a putative analytic expansion of astrophysically motivated waveforms. All ad-hoc waveforms have 1 second, half-Hann taper windows applied to the start and end and are sampled at 4096 Hz. Because they are all normalised to an amplitude of $10^{-20} \text{ Hz}^{-0.5}$, ad-hoc waveforms are useful for understanding the pipeline capabilities in a neutral context, as similar sensitivities are expected from waveforms with similar properties like peak frequency, bandwidth and duration [11].

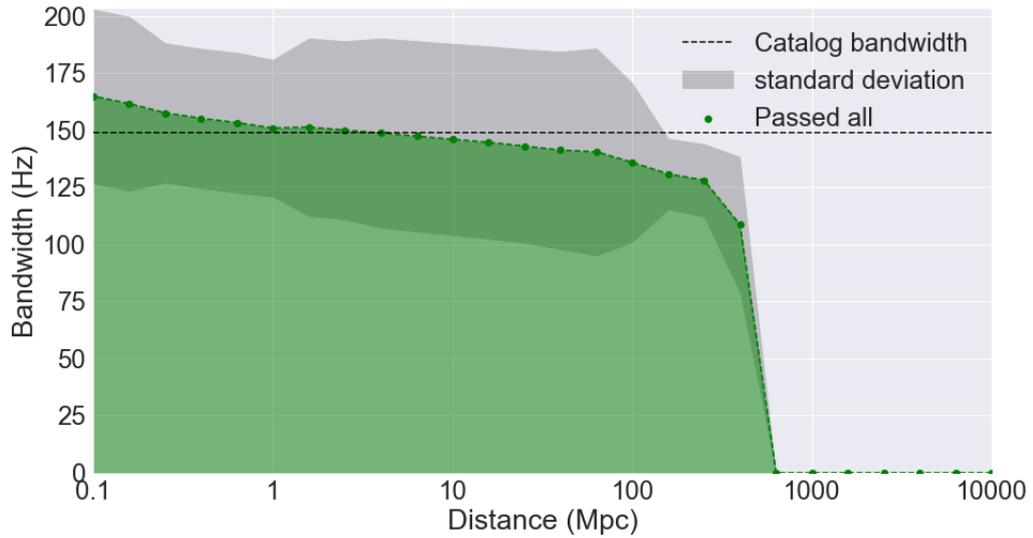


Figure 3.25: Recovered bandwidth for ADI-E: evolution of mean recovered bandwidth per distance from source. The shaded region corresponds to 1-sigma standard deviation. The dashed line shows the exact bandwidth of the waveform model. The waveform bandwidth is likely to be recovered accurately starting from 9 Mpc. The effects of Fourier leakage are observed up to 5 Mpc and artificially inflate the bandwidth recovery by saturating adjacent pixels, subsequently they are all clustered together.

3.7.2.1 Simple Models

The monochromatic, linear and quadratic waveforms are described by

$$(3.3) \quad x = A \sin(2\pi f(t)t).$$

Table 3.7 summarises the individual waveform properties:

Monochromatic waveforms like Mono-A (figure 3.32) and Mono-C (figure 3.33) are signals oscillating at a single frequency that does not change over time. As expected due to the data-cleaning process specifically targeting monochromatic-like signals, the pipeline is not well suited to detect signals with time-independent frequency.

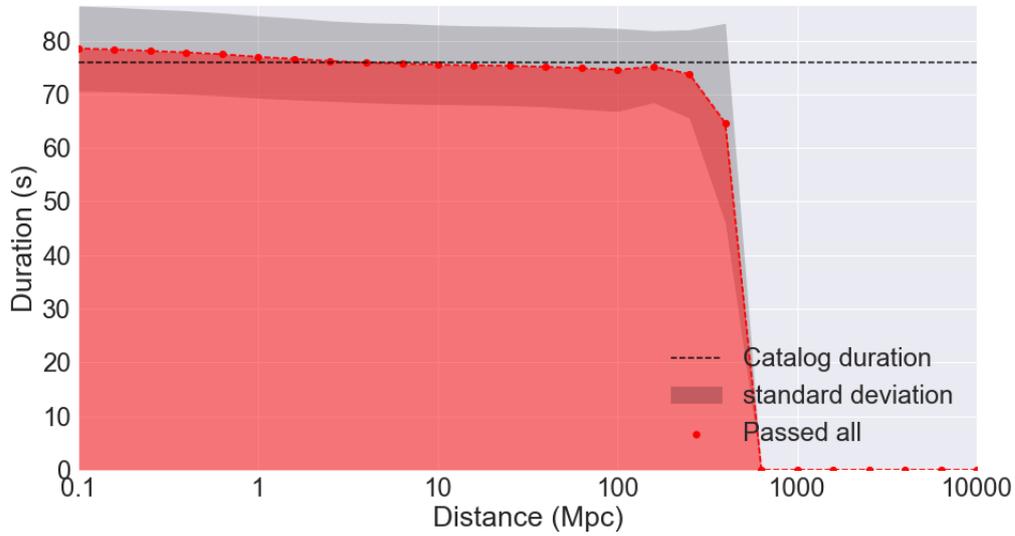


Figure 3.26: Recovered duration for ADI-E: evolution of mean recovered duration per distance from source. The shaded region corresponds to 1-sigma standard deviation. The dashed line shows the exact duration of the waveform model. The waveform duration is accurately measured from 5 Mpc onward.

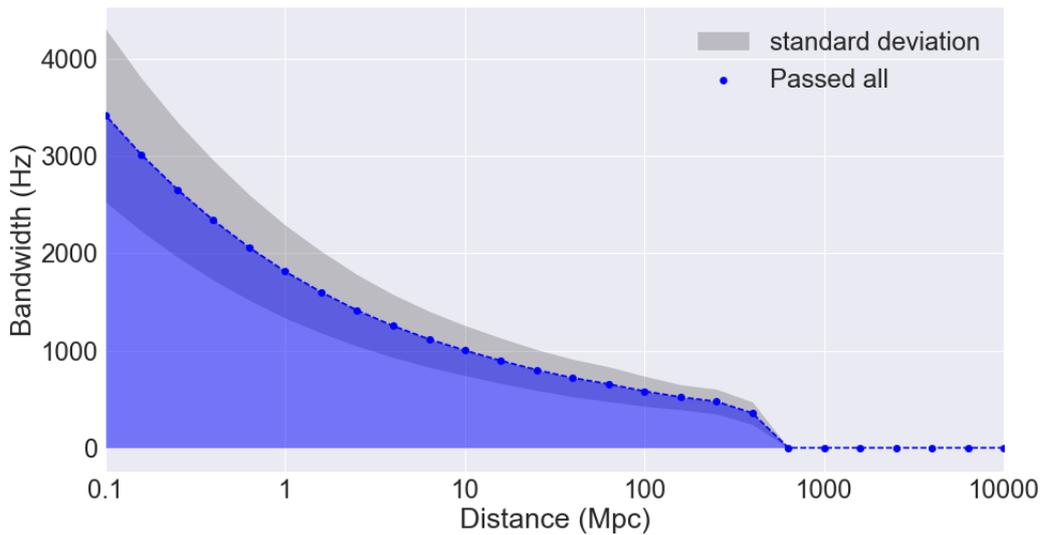


Figure 3.27: Mean number of pixels per trigger for each injection scale for ADI-E. The shaded region corresponds to a 1-sigma standard deviation.

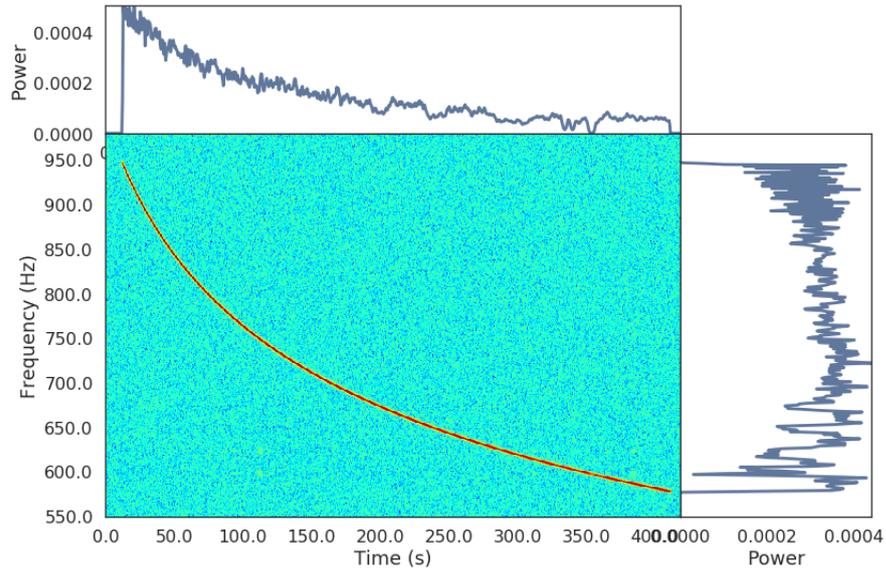


Figure 3.28: Energy spread of Magnetar-F in time and frequency. The waveform energy is decreasing over time with maximum at the beginning of the waveform. The spread of energy in the frequency domain is relatively constant, with a faster decrease towards the end of the waveform.

Waveform	Duration (s)	Start Freq. (Hz)	End Freq. (Hz)	Optimal SNR*
Mono-A	150	90	90	341
Mono-C	250	405	405	418
Line-A	150	50	200	915
Line-B	100	900	700	636
Quad-A	30	50	200	687
Quad-B	70	500	600	687

Table 3.7: Properties of interest for simple ad-hoc waveforms.

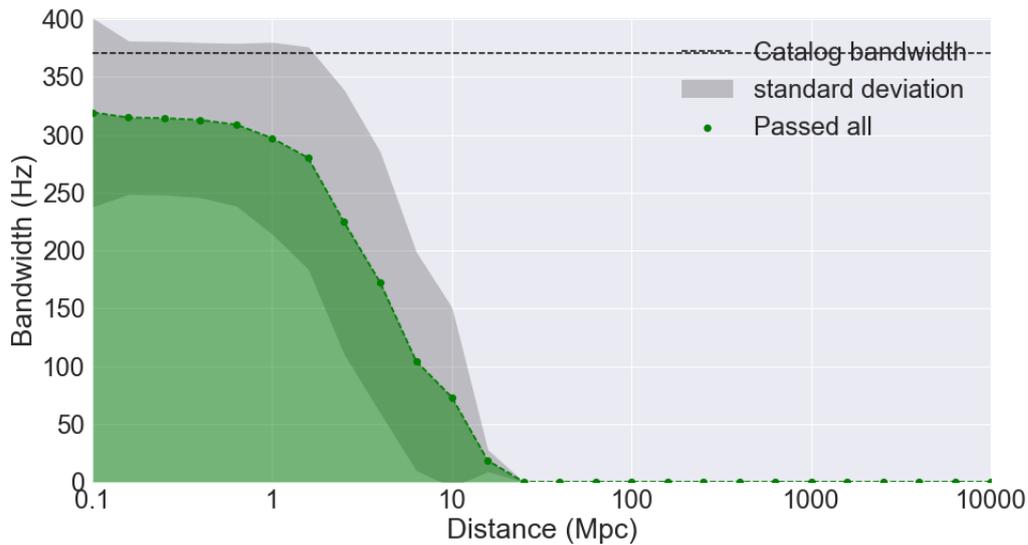


Figure 3.29: Recovered bandwidth for Magnetar-F: evolution of mean recovered bandwidth per distance from source. The shaded region corresponds to a 1-sigma standard deviation. The dashed line shows the exact bandwidth of the waveform model. The waveform bandwidth is never fully recovered and falls exponentially after 1 Mpc.

In cases where monochromatic-like waveforms do not cover the full 512 s block time, the remaining time-frequency bins will be over-whitened (see for example Figures 3.32 and 3.33)⁵ and will appear as pixels of lower brightness. The over-whitening of a particular frequency band effectively raises the threshold for subsequent pixels in the said band to be selected in the clustering step.

Line-A (Figure 3.34) and Line-B are signals whose frequencies increase or decrease linearly over time. Because of their linear increase in energy over time and frequency (Figure 3.35), they can be used to characterise the effect of the whitening filter and the efficiency of the clustering algorithm at low amplitude. However, the

⁵We note that this over-whitening effect extends to all waveforms with a monochromatic component. The late-time portion of the ADI family waveforms, for example, are affected as well. See Chapter 4 for improvements suggestions in the next observing run.

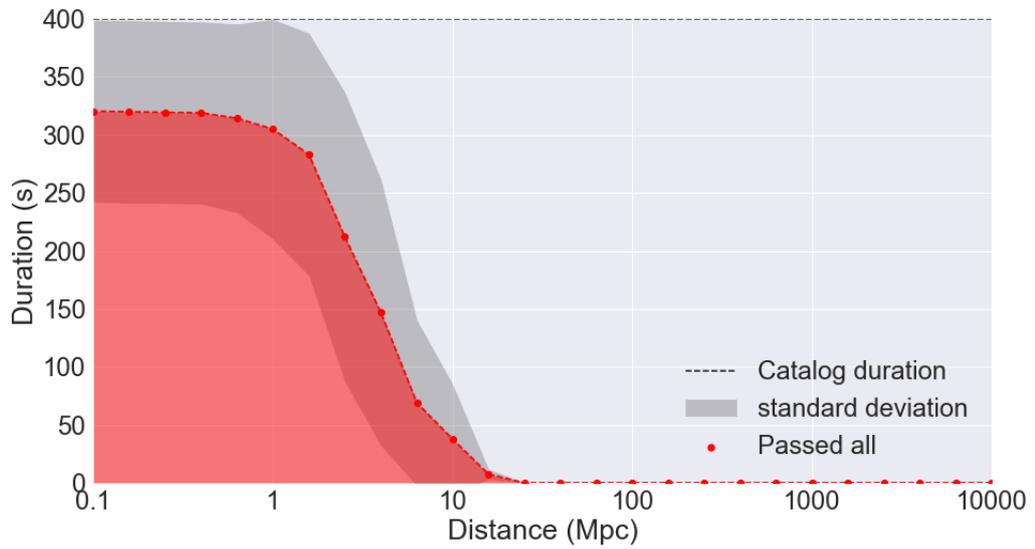


Figure 3.30: Recovered duration for Magnetar-F: evolution of mean recovered duration per distance from source. The shaded region corresponds to 1-sigma standard deviation. The dashed line shows the exact duration of the waveform model. The waveform duration is never fully recovered and fall exponentially after 1 Mpc

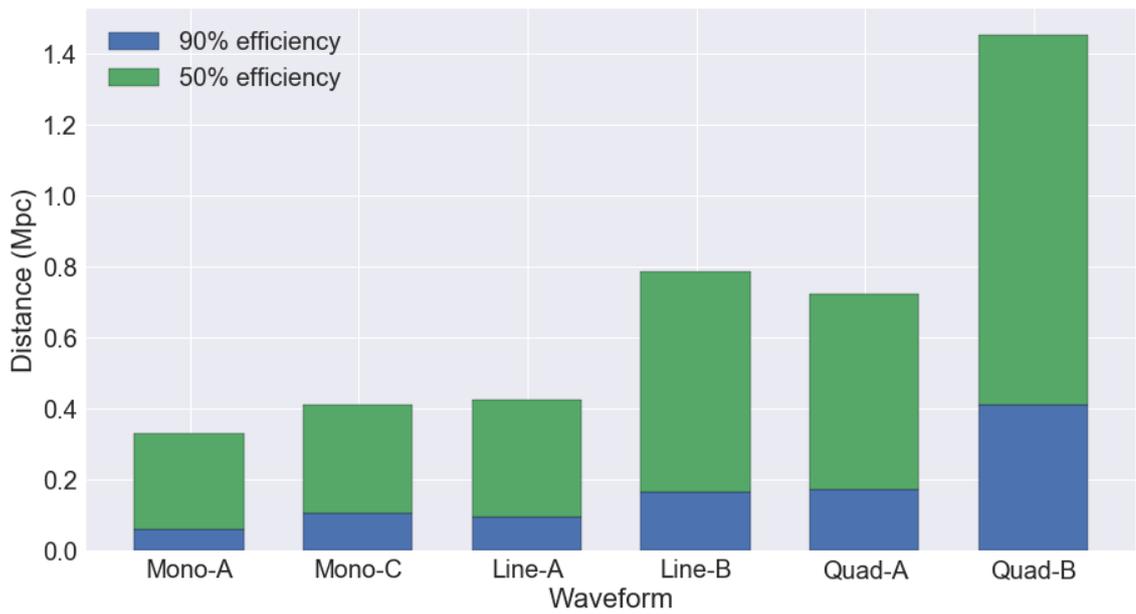


Figure 3.31: Effective distance for the simple ad-hoc waveforms during the O1 period at a FAR of 1 in ~ 50 years.

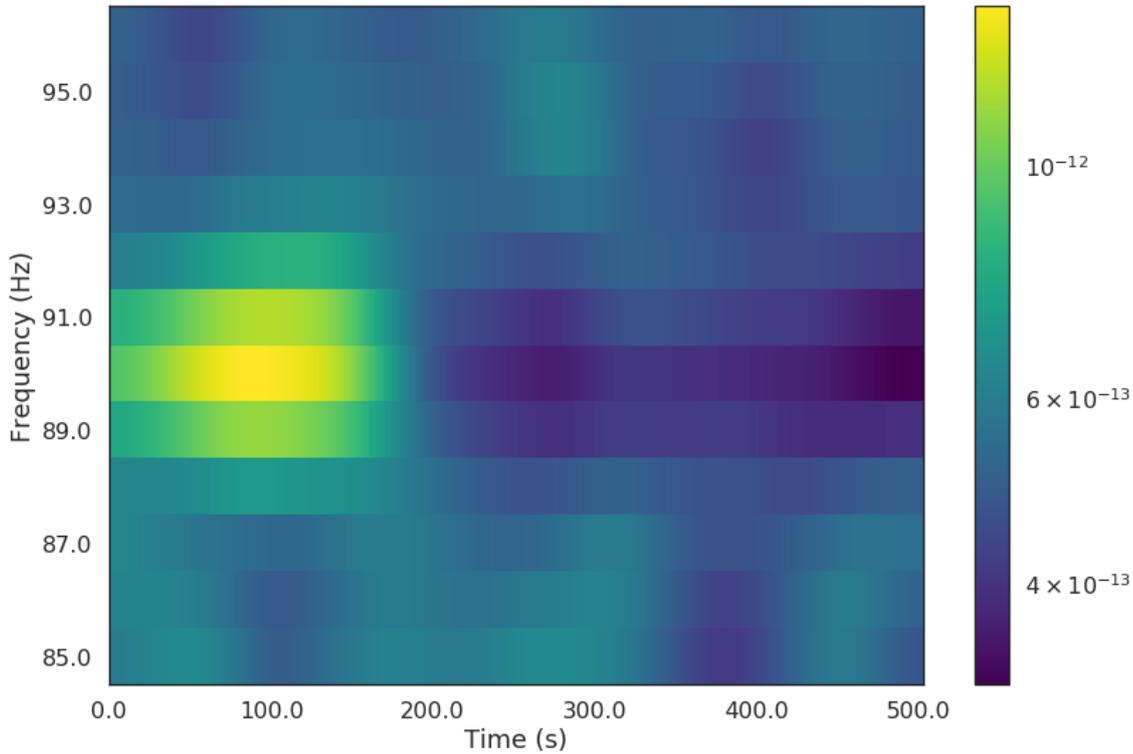


Figure 3.32: Low amplitude software injection of MONO-A waveform. The waveform is visible at ~ 90 Hz, from 0 to 150s. The effect of the zero-phase linear predictor error filter is visible from 150 s onward, where a clear line of less-than-average power is present at the same frequency. The effect of the filter has been enhanced using a gaussian average window over 51 seconds.

smart clustering step discards the useful information when combining bounding boxes neighbour clusters.

Because of its higher frequency and the frequency bias of the alpha cut, Line-B is recovered nearly twice the distance, or half the amplitude, than Line-A at 0.7 Mpc. No breaking of the waveform occurs, indicating that lower frequency triggers are more likely to be broken up into segments.

Quad-A (Figure 3.36) and Quad-B (Figure 3.37) are more likely to be successfully recovered than the simple ad-hoc models due to their chirping properties - the smart

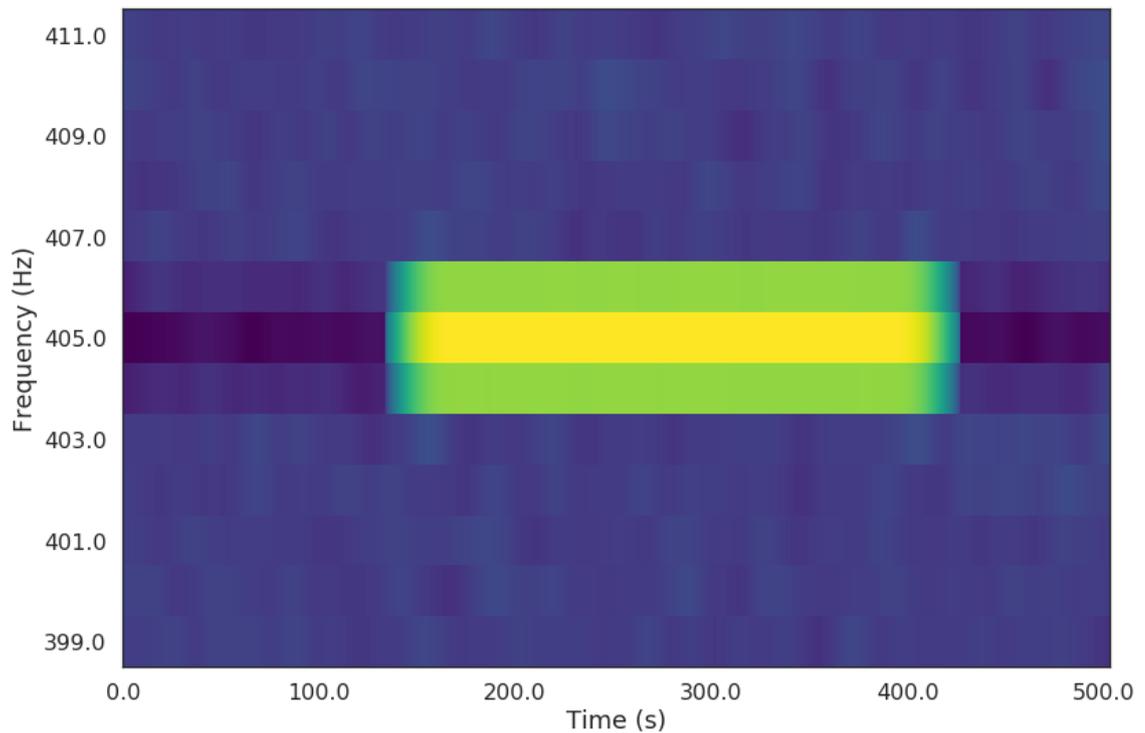


Figure 3.33: High amplitude software injection of MONO-C waveform. The waveform can be seen at 405 Hz, between 150 to 400 s. The effect of the zero-phase linear predictor error filters are seen up to 150 s and after 400 s onward, where a clear line of less-than-average power is present at the same frequency. The effect of the filter has been enhanced using a gaussian average window over 11 seconds

clustering is favouring a small number of pixels in a large time-frequency bounding box. Similarly to the waveforms discussed so far, X-SphRad is more sensitive to Quad-B due of the efficiency bias toward high-frequency. The bias is slightly offset by the longer duration of the Quad-b waveforms compared to Quad-A waveforms.

3.7.2.2 Gaussian Family

The sine-Gaussian waveforms are simple unipolar monochromatic signals described by eq. 3.4 multiplied by a Gaussian envelope.

$$(3.4) \quad h = e^{-t^2/\tau^2}.$$

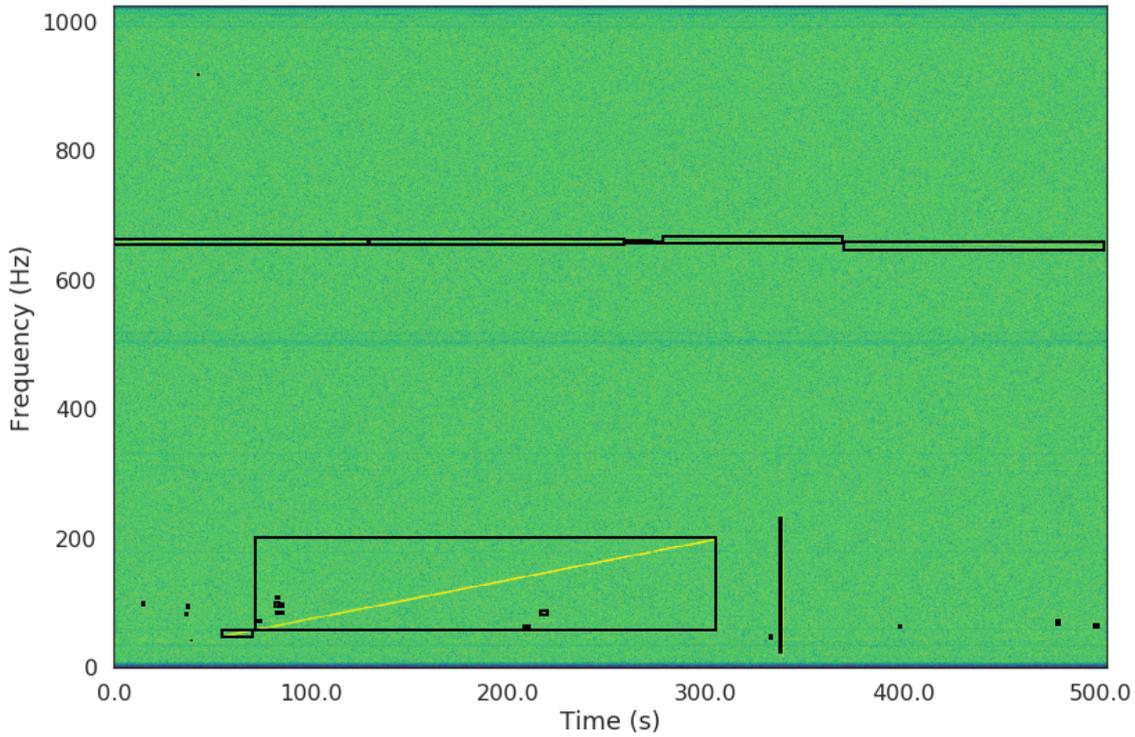


Figure 3.34: High amplitude software injection of Line-A waveform. Potential GW candidates selected by the pipeline are denoted by a black rectangle corresponding to the event bounding box. The waveform can be seen at ~ 50 Hz to 200 Hz, from 60 to 300 s. Despite the high amplitude of the injection, the pipeline split the waveform at 80 s, and clusters glitches (at 90 and 200s). The wandering line glitch at 630 Hz is also selected as potential GW candidate, as well as a wide-band glitch at 340s.

As they are similar to the Mono family, they are penalised by the whitening filter in the same way. Figure 3.38 shows the effect of the whitening filter; at lower amplitudes the whitening filter effectively removes the sine-Gaussian (top), and corrupts the frequency bin they vibrate at for following trigger events in the block at higher amplitude (bottom). For more moderate amplitudes, only the central part of sine-Gaussian is recovered as the start and end of the waveform is comparatively weaker than the central part (middle).

The White Noise Burst (WNB) waveforms are white noise band-passed signals

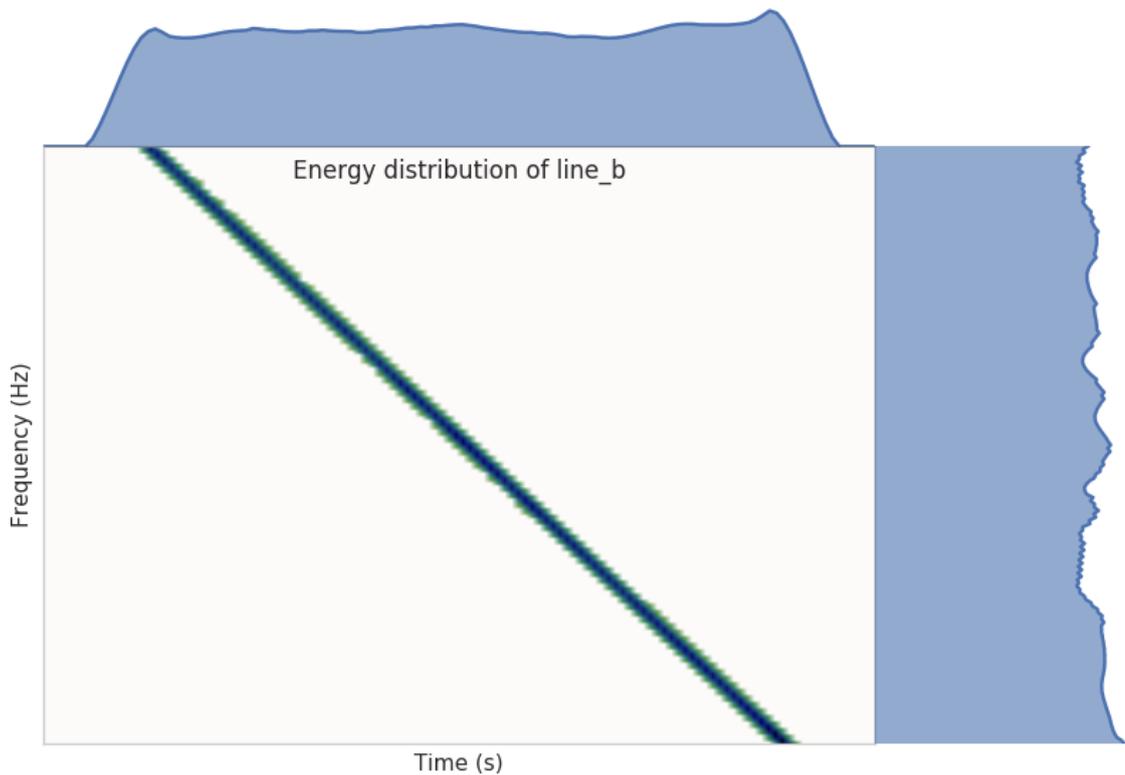


Figure 3.35: Energy spread of Line-B in time and frequency. The waveform energy over time is relatively constant, with a hann window taper at the beginning and end. The spread of energy in the frequency domain shows an inverse chirp with power concentrated in the end-life of the waveform.

with a Gaussian time envelope. As they closely resemble typical instrumental glitches like the blip glitch (Figure 3.39), the WNB waveforms are useful to understand the impact of long wide-band glitches in the trigger selection process. A high and medium amplitude WNB are split into several follow-up candidate events. The effect is more pronounced at medium amplitude, where the WNB is split into more than 20 follow-up events out of the 30 per time-frequency map.

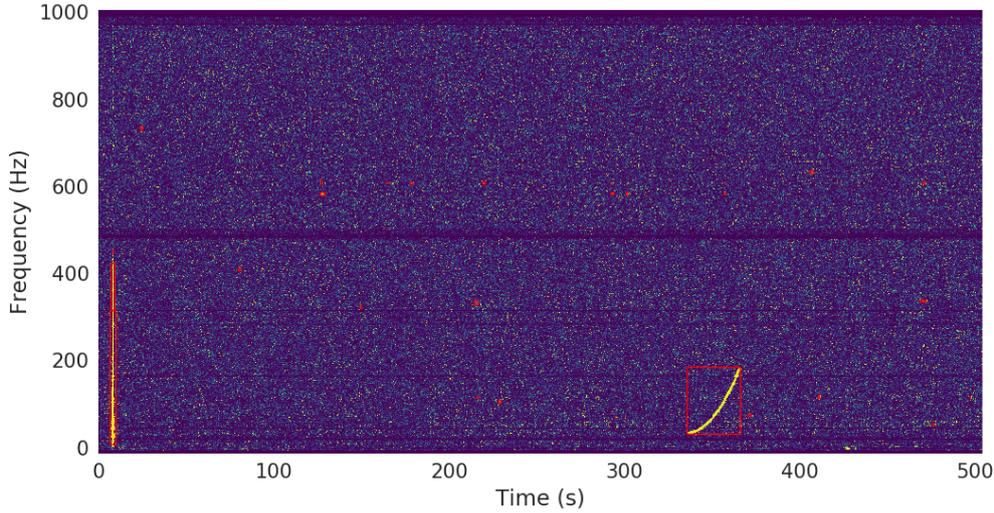


Figure 3.36: High amplitude software injection of the Quad-A waveform. The waveform can be seen starting from 330 s to 360 s, with the bounding boxes selected by the pipeline in red. A loud line glitch is visible at ~ 10 Hz.

3.8 Upper Limits and Pipeline Comparison

This section is adapted from [22]. The search reports a few low-significance candidate events that are consistent with background noise. As a consequence, no significant events were observed and we use the loudest event procedure [38] in order to set upper limits on the rate of long duration GW transient signals at 90% confidence following Poisson' statistics, updating those set in [23]. The uncertainty from calibration errors is folded in the calculation.

For a Poisson distribution and an observing time $T = 49$ days, the number of

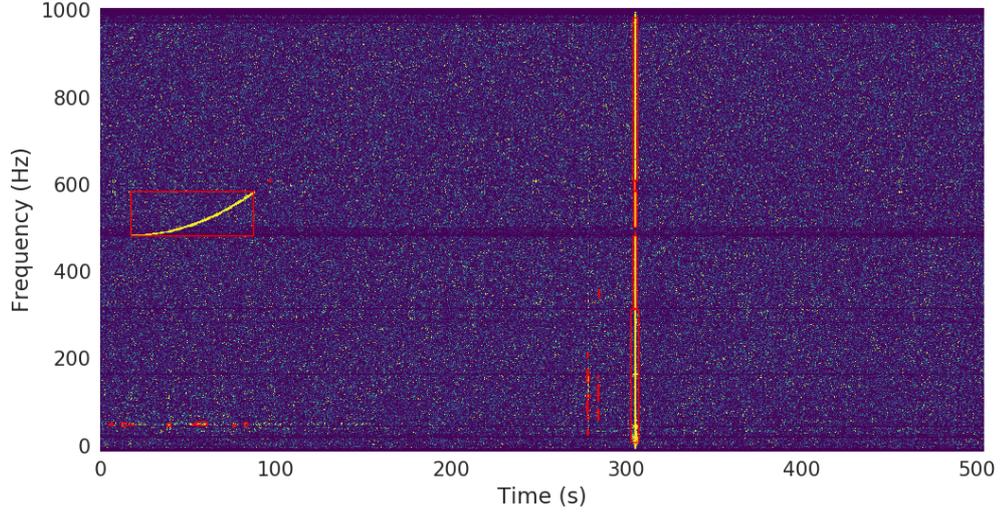


Figure 3.37: High amplitude software injection of the Quad-B waveform. The waveform can be seen starting from 30 s to 90 s, with the bounding boxes selected by the pipeline in red. A loud line glitch is visible at ~ 10 Hz.

expected event is

$$(3.5) \quad n_{events} = e_0 \lambda * T$$

where e_0 is the efficiency at a particular distance.

The probability of seeing n GW events is

$$(3.6) \quad p(n|e_0 \lambda * T) = \frac{(e_0 \lambda * T)^n}{n!} e^{-e_0 \lambda * T}$$

and so the probability of seeing 0 event is

$$(3.7) \quad p(0|e_0 \lambda * T) = e^{-e_0 \lambda * T}$$

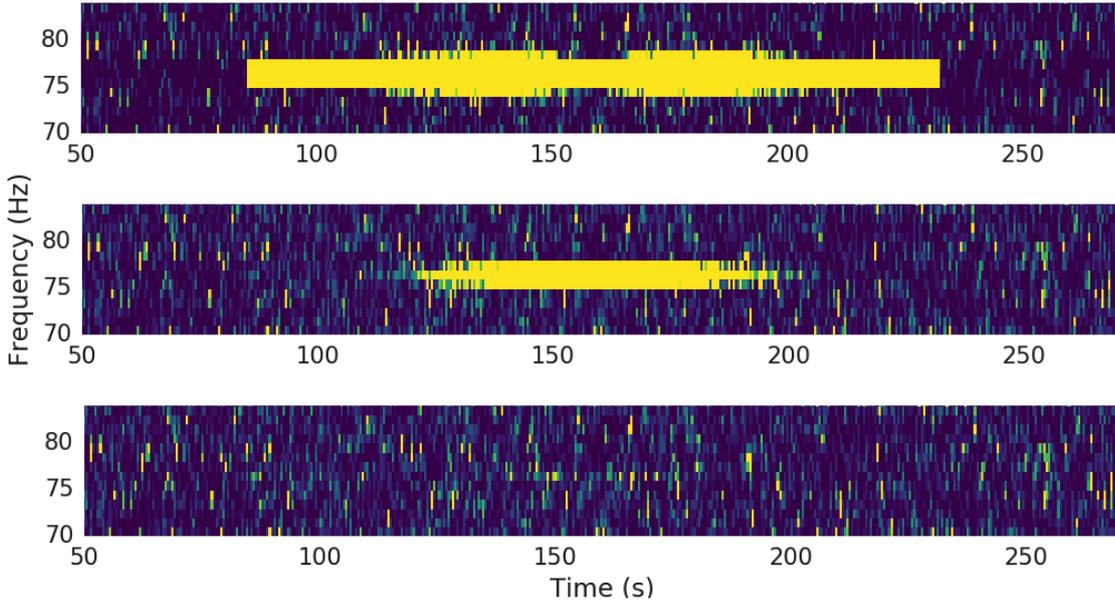


Figure 3.38: Comparison of the whitening filter effect on a sine-Gaussian waveform. The overly loud signal (top) will corrupt the filter and downweight the values of the other pixels in the frequency bins. For moderate amplitude, the start and end of the waveform are discarded due to the Gaussian envelope (mid). At lower amplitude (bottom) the whitening filter effectively removes the sine-Gaussian.

with thus a probability of observing an event of

$$(3.8) \quad p(\text{not } 0 | e_0 \lambda * T) = 1 - e^{-e_0 \lambda * T}$$

Setting a 90% confidence limit threshold gives $p = 0.1$ which leads to

$$(3.9) \quad p(\text{not } 0 | e_0 \lambda * T) = 1 - e^{-e_0 \lambda_{90\%} * T}$$

We then have

$$(3.10) \quad \frac{1}{10} = e^{-e_0 \lambda * T} \Rightarrow \lambda_{90\%} = \frac{2.3}{-e_0 * T}$$

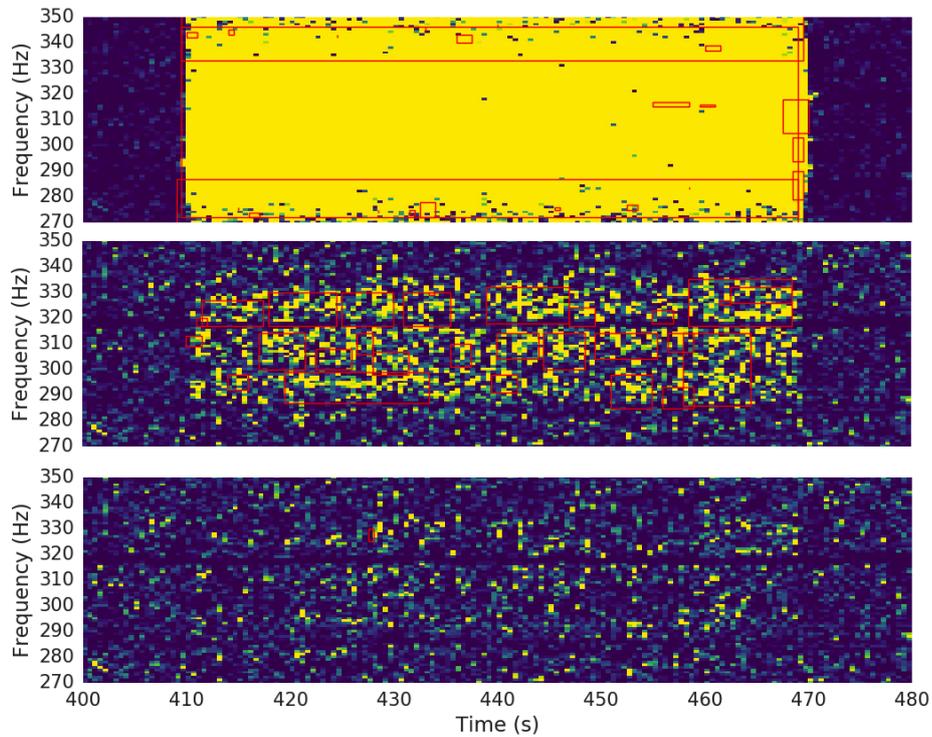


Figure 3.39: Comparison of the recovery of a WNB waveform. The overly loud signal (top) resembles a loud wide-band glitch, which corrupts the filter and downweights the values of the other pixels in the covered frequency bins. For moderate amplitude (mid), the waveform has the same time-frequency properties as blip glitches. Even at lower amplitude (bottom), the clustering algorithm follows broken fragments of the waveform.

Table 3.8: 90% confidence upper-limits for the astrophysical waveforms.

Waveform	90% confidence limit (Mpc ³)
ADI-A	0.0021
ADI-B	0.0001
ADI-C	0.0014
ADI-D	0.0226
ADI-E	0.0007
Magnetar F	11.766
Magnetar G	99.982

Table 3.8 reports the 90% confidence limit threshold.

ENHANCEMENT TO LONG-DURATION BURST SEARCHES

4.1 Introduction

Unmodelled searches typically look for an excess of power in a spectrogram to flag GW candidate events [30]. Clustering algorithms assign [63] pixels into groups that are further tested against consistency cuts and vetoes to determine the likelihood of the presence of a signal. Most of the X-pipeline processing time is spent gathering data, correlating it and generating a time-frequency map. A comparatively small fraction is spent localising and analysing GW candidate clusters. The clustering phase is relatively unexplored and unoptimised, for example LIGO dcc search with the term 'clustering' returns only a handful of relevant results (referencing notably [42, 45, 5]). However, it is the one of the principal factors limiting the X-SphRad O1 search sensitivity.

In all search pipelines, there is a trade-off between sensitivity and computational time. Analyses need to be sensitive and have a non-zero potential for detection, while providing results in a timely manner. For searches such as the one presented in Chapter 3 the background estimation is by far the most computationally costly procedure, with ~ 50 years of data to be analysed in a few months. Adding a few extra seconds to any steps of the pipeline has repercussion that can be counted in days of extra analysis time. In this context, the value of any increase in amount of time spent finding events of interest on the time-frequency map depends on the resulting increase in sensitivity [24].

As seen in the previous chapter, glitches significantly hinder the sensitivity of long-duration GW transient searches including X-SphRad. As the number of follow-up candidate events on a time-frequency map is limited, X-SphRad ranks then selects top priority clusters based on their sky energy. This naturally favours loudest events on the time-frequency map. As loud glitches typically occur at least once during a 512s interval, most if not all of the follow-up triggers selected by the pipeline thus originate from transient noise. When looking for an excess power for minute-long signals, relatively low amplitude signals would be missed.

We propose a novel method for GW clustering on time-frequency maps that uses properties beyond the amplitude of the potential signal. By using an edge detection algorithm, we quantify each pixel in the time-frequency map by similarity with its local neighbours and extract features of sharply changing intensity (or 'edge'). We test the method on a subset of injections from O1 and find that we detect quiet events consistently better than currently implemented algorithm, with an increase in observable volume by nearly an order of magnitude. The method assumes only that the potential signal is represented by at least few pixels on a time-frequency map

with higher amplitudes than its local neighbours in the time-frequency plane.

4.2 Clustering the time-frequency map

The simplest form of clustering algorithms can be applied like an excess power search, and generate clusters based only on the pixel values. The algorithm first selects a threshold value, and discards all pixels in the time-frequency map with lower values. The pixels sharing a side or corner are kept and flagged as a single cluster. The method is fast, simple to implement and work well in environments with low-noise and high SNR signals. To make it more robust, the algorithm can also apply a second lower threshold (sometimes called "core or halo clustering"), and treat the first threshold as seed from which branches grow up to reaching a pixel with a value lower than that of the second threshold. This is appropriate for events where at least one pixel of the cluster has high amplitude, and some of the values of the cluster pixels are below the initial threshold. Branches can also be allowed to "jump" (or "leapfrog") a certain distance and continue growing as normal.

In environments with non-Gaussian non-stationary noise of high amplitude, and target signals with low SNR such as in the case of the LIGO detectors, the method are not necessarily optimal. The first threshold will discard interesting events if set too high, and will aggregate important quantity of noise if set too low. The second threshold allows for weaker features to be extracted, but is similarly easily corrupted and can still miss pixels with significant signal energy. Adding a jump distance is not necessarily optimal neither, as it can similarly lead to over-selection and event data corruption. The clustering algorithm described in section 2.6 tries to mitigate the effects of the non-stationary background noise by dividing the parameter-space into

regions with individual thresholds, but still lose signals due to glitches in individual regions (Figure 4.1). Pixels available for event reconstruction after thresholding are limited and sparse.

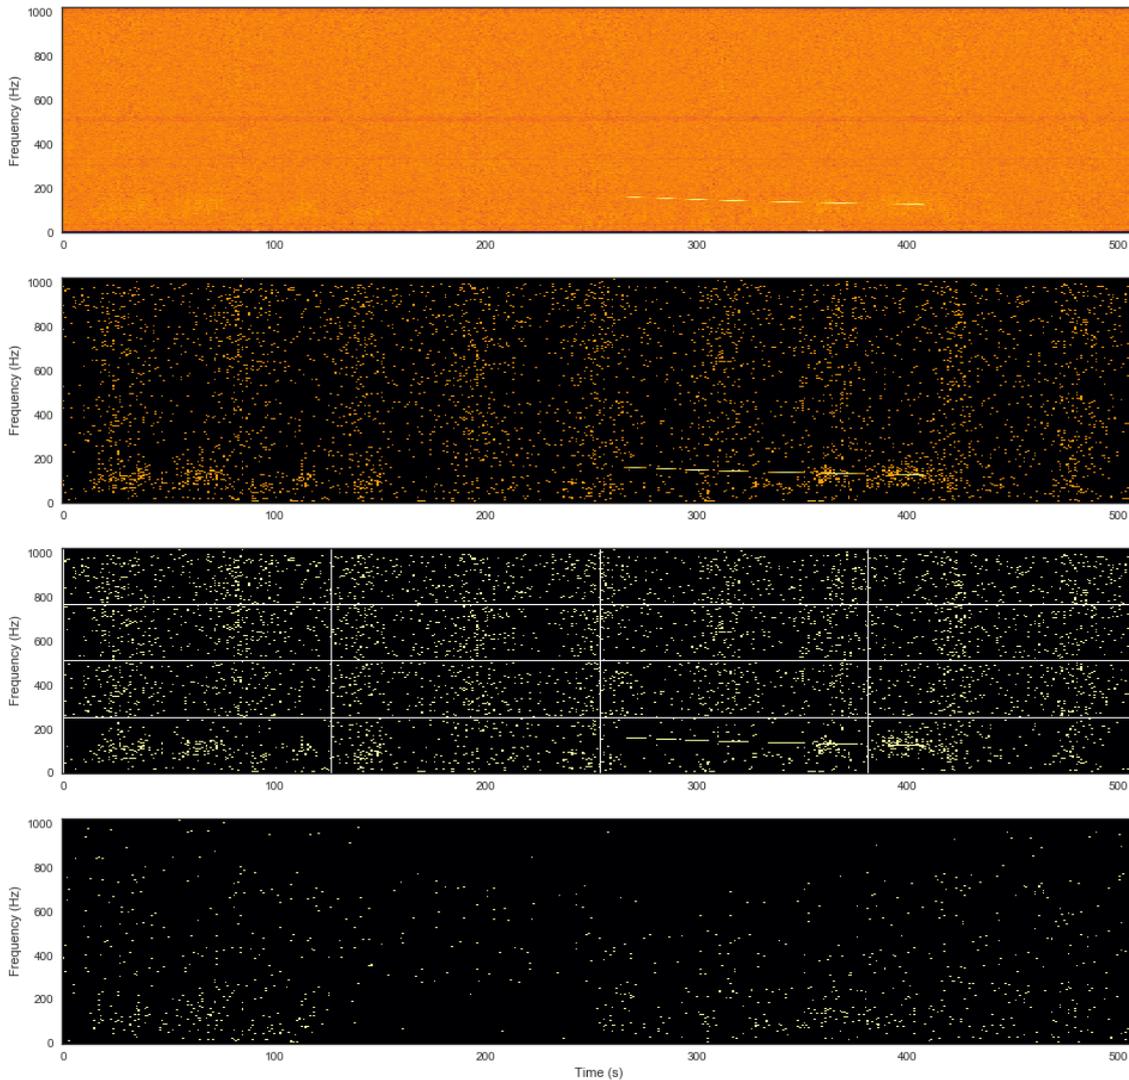


Figure 4.1: Comparative time-frequency maps showing the effect of clustering in X-SphRad, with and without region-based thresholding. A simulated signal of moderate amplitude from 260s to 430s. (top) unclustered time-frequency map. (2nd) Pixels surviving loudest percentile threshold. (3rd) Pixels surviving the region-based threshold. (bottom) Logical XOR between the two methods. Adding regions shows only a marginal difference in pixel selection.

Other long-duration GW searches ([71, 46]) have developed alternative solutions to extract information from noisy environments. They generate masks of the desired shape, and maximise the integrated power along the mask. For example, coherent WaveBurst (cWB) [71] select GW candidate events by performing a convolution with a kernel shape boosting chirp-up or chirp-down. However, the kernel biases against waveforms that are not chirp up or chirp down by construction. It also loses the potential to discover the most interesting GW events that are not expected. STAMP lonetrack [46] generates series of Bezier curves under the assumptions that GW emissions appear as curvy tracks on the time-frequency map, which is computationally intensive. Both are looking for general excess of power in the data.

In excess power clustering algorithm the critical assumption is that the candidate GW event will be represented on the time-frequency map by at least one pixel of power higher than or comparable to a high percentage of the background. This is most often not the case for long-duration analyses with block time in the minute range, as the LIGO detectors exhibit frequent non-stationary noise fluctuations of high amplitude. In this light, we argue that excess power clustering methods are not the best suited method for minute-long GW transient searches. For long duration signals, by definition the signal energy will be spread out over many pixels, so looking for individual very-high amplitude pixels is unlikely to be the optimal strategy.

We can define an optimal algorithm in the context of long-duration GW clustering as having the following ideal properties. The perfect clustering should not be hypersensitive to glitches, but should nonetheless not be oblivious to them and it should flag them as they exhibit behaviours similar to real signals. The distinction between GW event and GW candidate should be left to follow-up steps. The perfect

clustering method should be truly unmodelled and not target specific time-frequency properties of current best approximate waveforms. It should be fast and efficient enough to run several-sigma background in reasonable time and increase the pipeline sensitivity.

With the proposed method, we are making the assumptions that long duration GWs resemble tracks in the time-frequency space, based on long duration signal models looking like lines (or tracks) on the time-frequency map. However, broadband long duration GW signals would have their power spread over more pixels and be inherently less detectable. The assumption is thus minimal. Finding such tracks is an important part of the field of computer vision and algorithms are used extensively in radiology (see [73, 59, 48] for example). The same techniques are used for image enhancement and artistic effect in image processing [106]. Image processing techniques applied in this context allow for efficient algorithms that have the advantages of being local, adaptive, and fast.

4.3 Edge detection

Edge detection is a common problem in image analysis [104]. Edges are often used to bound objects [89] and have applications in image segmentation [90], feature extractions [53, 28] and object identification [88, 78, 68]. The technique is also used to enhance image features [82]. The most common algorithms of edge detection are Sobel [49], Canny [41], and Laplacian [64]. Edges are considered as high frequency features so techniques are akin to high-pass filtering of the time-frequency map and are essentially a convolution of a kernel with the original input, so intrinsically local within the chosen kernel size.

The Sobel filter is a convolution with a kernel approximating the directional first derivative. For example, a 3×3 kernel is

$$(4.1) \quad \text{Sobel : Convolve with } G_x = \begin{bmatrix} +1 & 0 & -1 \\ +2 & 0 & -2 \\ +1 & 0 & -1 \end{bmatrix}.$$

capturing the gradient changes in the x direction. The kernel can be rotated by 90 deg to target the y direction then addition in quadrature for overall power. Due to the weights assigned in the kernel, the method only finds edges in a specified direction. The simple convolution can be computed quickly but the method is very sensitive to noise fluctuations. Other kernels are available that are more isotropic and larger [36]

The Canny filter [41] is an enhancement to the Sobel filter, building up on it with several extra steps. The first step is the blurring of the initial image to minimise the noise fluctuations before convolving on a Sobel kernel (or alternative kernel) in the x,y and the two diagonal directions. To compute the final value for each pixel, the four derivatives are added in quadrature. The principle behind the initial blurring is that the noise is a high frequency component of the image, so the image is low-pass filtered. However, sharp edges are also high frequency components, and so the blurring kernel size needs to be carefully chosen to balance noise reduction to edge detection sensitivity. As edges are widened by the blurring, the pixels that are not local maximum are discarded to thin the edges. A pair of threshold are used to identify edges, and weak edges that do not touch strong edges are discarded. The operation has a few more steps than Sobel, and is thus computationally more intensive. Enhancements to the technique [47] speed up the computation, but the

method is best suited for mid-SNR environment. The double thresholds discard most of the edges caused by spread noise, but can distort target signals by flagging noise physically connected (Figure 4.2).

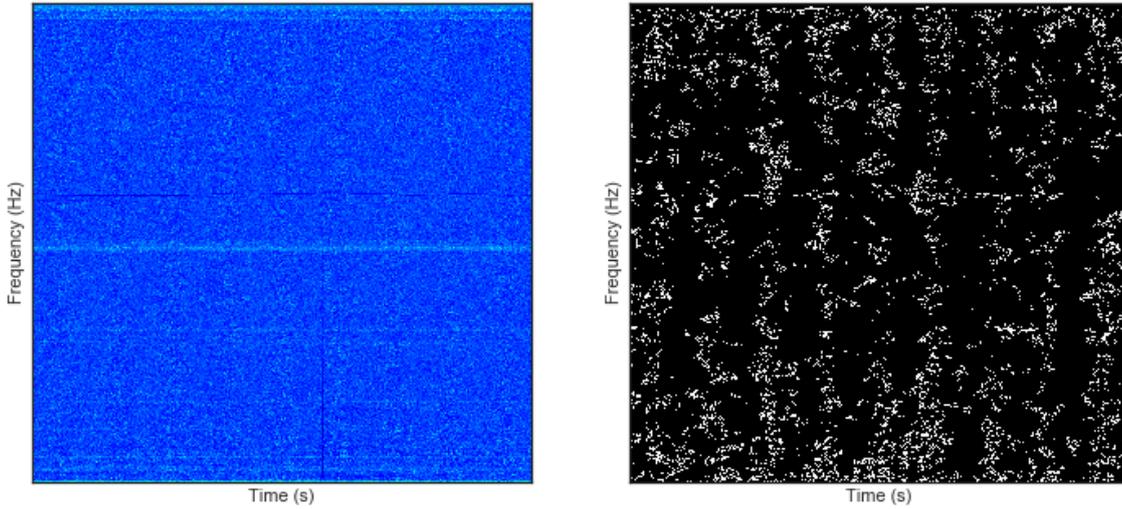


Figure 4.2: Effect of the Canny algorithm on a time-frequency map. (left) Original time-frequency map (right) Canny filtered. The canny-filtered time-frequency map does not recover fully the high-amplitude 630 Hz wandering line glitch, and clusters it with additional surrounding noise. Other features like the short-lived broadband glitch at the centre of the map are lost in the accumulation of noise.

Other edge detection methods such as the Laplacian use a second-order derivative of the intensity instead [64]. This way, it is not the change of gradient that is estimated, but the rate at which the gradient changes with zero-crossings indicating local maxima in the gradient. The laplacian operator in two dimension is

$$(4.2) \quad \nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

which can be approximated on a discrete grid by a convolution of the kernels of the

form

$$(4.3) \quad \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}, \text{ or } \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix}.$$

Second-derivatives are very sensitive to noise however, and an initial blurring pre-processing step is needed for the method to be effective in noisy environments, as with Canny's . The convolution operation is associative, so the computation can be sped up by convolving the Gaussian filter and Laplacian filter together, then convolve the result with the data input, requiring only one pass through the data.

The most interesting difference between the Laplacian and Canny's is the evaluation of the rate of change in gradient, instead of the gradient itself. Using similar techniques for GW detection on time-frequency map allows to threshold on the rate of change with respect to the local background, instead of pure signal-to-noise ratio. While loud glitches are limited in time-frequency and exhibit high rate of gradient change compared to the background. Our target GW signals are more localised as thin lines in time-frequency regions. We can reduce the sensitivity to noise by extending the kernel size, and sporadic pixel of noise excess will be averaged together with the potential signals. However, the method allows only for control of the upper-limit of the low-pass filter from the initial blurring step.

4.4 Difference of Gaussians

The Laplacian can be approximated by a difference of Gaussians (DoG) [75]. By convolving a Gaussian kernel of size m with the input, then subtracting a convolution with a second Gaussian kernel of size $n > m$, the difference between the two

approximates well the laplacian and effectively acts as a band-pass filter. The Gaussian kernel sizes control the bandwidth, as the subtraction of one convolution with the other will only preserve the spatial information within the range of frequencies preserved in the two blurred images [33]. Effectively, increasing the larger radius n lowers the threshold for recognising the edges and increasing the smaller radius m will thicken them [34]. The resulting filter is isotropic, so the algorithm is fast as it need only one pass in the data. It can even be speed up by approximating a Gaussian with successive box blurring with uniform boxes, see Figure 4.3 [94].

Figure 4.5 shows the effect of the DoG on a standard time-frequency map containing a relatively loud simulated signal. Increasing the size of the larger kernel n translates in a loss of details and generally larger "blobs" of overall low gradient. Increasing the size of the smaller kernel m effectively clusters more pixels around areas with high rate of gradient change. Selecting more pixels around areas of interest could potentially be useful for background rejection techniques. The Laplacian is best approximated [33] with the ratio

$$(4.4) \quad \frac{m}{n} = 1.6.$$

4.5 Preliminary results

The current implementation of the DoG and its attached cluster selection routines is done in Python. Several of the steps described below were only taken to empirically produce thresholded time-frequency map, on which cluster could easily be located using native python functions. These are not necessarily desirable in X-SphRad.

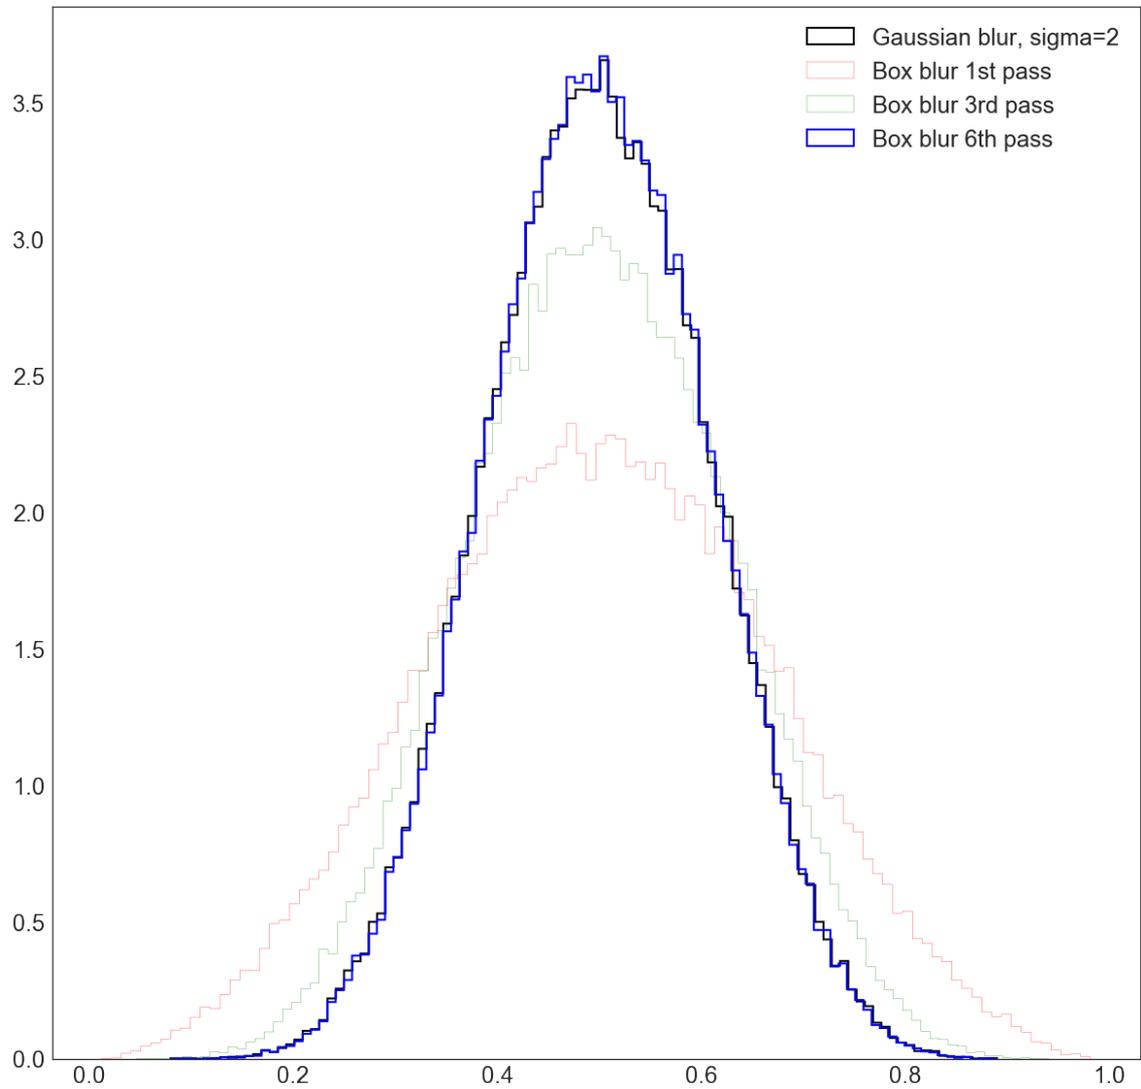


Figure 4.3: Histogram of values of a noise map after Gaussian blur and after box blur. The distribution of map values is essentially the same with both method.

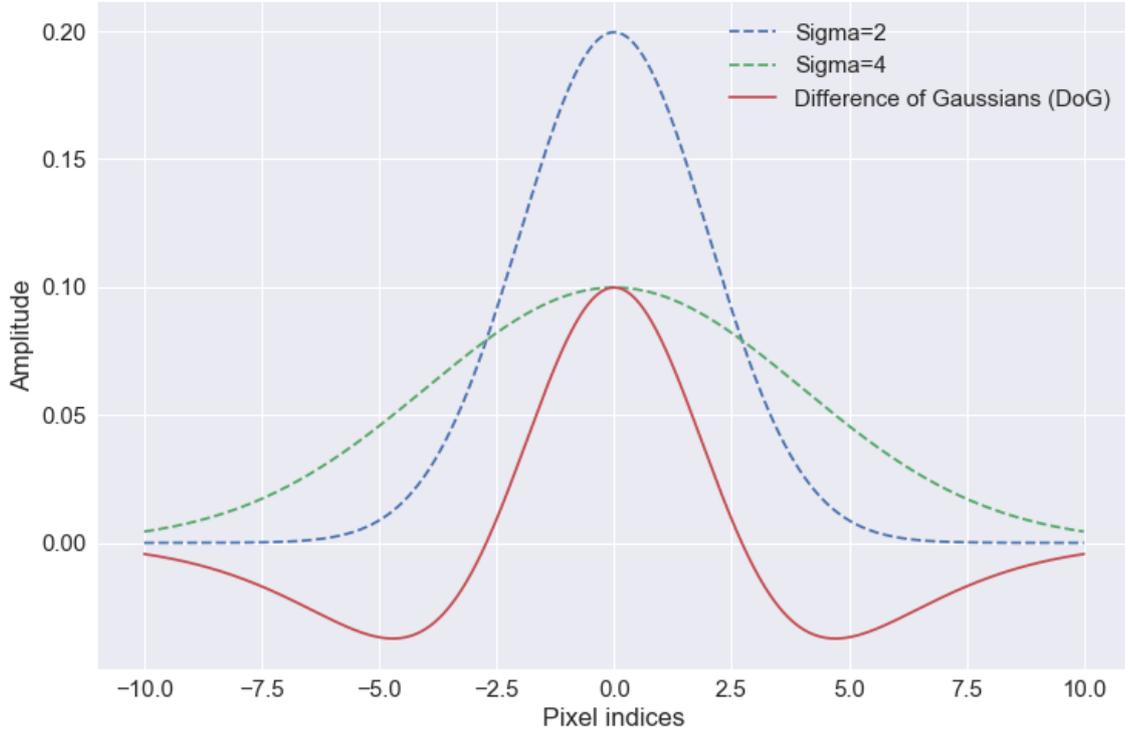


Figure 4.4: Difference of Gaussians on pixel weights, illustrated with value of 2 and 4 sigmas. The resulting weighting factors are commonly referred to as 'Mexican hat' and are used to enhance edges in image processing.

First, we load a time-frequency map generated by X-SphRad containing simulated signals. The map is rescaled by dividing by the minimum value in the map, then we take the logarithm of all the values. The logarithm step gives more details when observing the DoG algorithm visually and prevent overly loud glitches from saturating the colour bar on plots. The bands 0-60 Hz and 900-1000 Hz are reset to be equal to the minimum map value. Coherent noise in those frequencies has high amplitude and is relatively constant, which generates long and sharp edges that were consistently flagged as potential GW candidates. We apply the DoG with parameters $m = 7, n = 11$ (close to the 1.6 ratio to approximate the Laplacian) on a time-frequency map, then discard all pixels with negative value. We apply

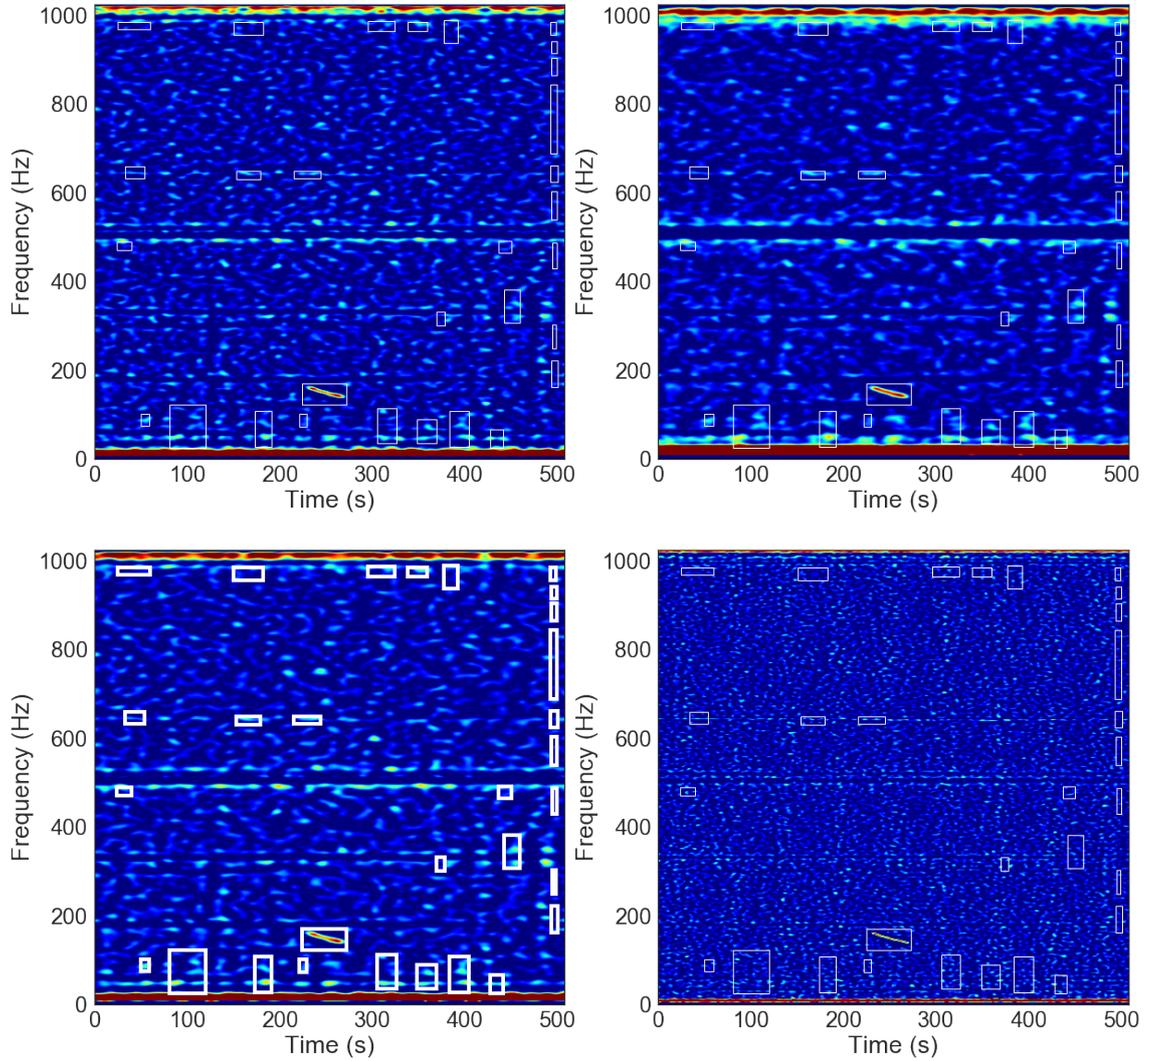


Figure 4.5: Effect of the DoG on a time-frequency map generated from data from the LIGO O1 period, with different kernel sizes and an added simulated signal. The differences due to kernel sizes are apparent in the excess of noise in the bottom and top 40 Hz of the time-frequency maps, and on the simulated signal. top-left) $m=2, n=7$ top-right) $m=2, n=12$ bottom-left) $m=5, n=7$ bottom-right) $m=2, n=3$. All time-frequency maps were thresholded by removing pixels with value lower than zero. For production mode, we recommend parameters with a ratio following Eq. 4.4

a Gaussian blurring with kernel size 3×3 to connect blobs that might have been split by the whitening filter or by noise, then threshold on the median value of the DoG map plus 2 standard deviations. We repeat the blurring thinning process with a second harder threshold of DoG map plus 5 standard deviations to have clear definition in the resulting cluster. We then find isolated islands of pixels and flag them. We rank all triggers formed of at least 10 pixels by their original sky statistic, weighted by their value on the DoG map. To compare fairly between the DoG and the clustering used by X-SphRad, we keep only the top 30 clusters with the largest weighted values. To be able to compare between the two methods, we generate a loud injection guaranteed to be detected and select pixels that are clearly from the simulation using the DoG. We then compare them to the pixels returned by the DoG and X-SphRad or moderately loud injections of the same signal.

We generated a set of ~ 80 injections for the ADI models D and E, using the same random distribution parameters and distance scales as per the O1 analysis. To be flagged as recovered, a cluster needs to overlap with at least 1 pixel of the simulated signal and have a sum weighted amplitude amongst the top 30 clusters of the time-frequency map. To allow the characterisation of the differences between the two methods, no further consistency cuts or thresholds were applied to the clusters including the loudest background event criterion used by X-SphRad. The full comparison including background rejection would require a complete analysis of the data (including background) with the DoG procedure, which is not feasible until the algorithm is implemented in X-SphRad.

Figure 4.7 and Figure 4.6 report the increase in candidate selection sensitivity for ADI-D and ADI-E. At low distances both methods recover almost all injections and the same number of pixels. A small fraction of the injections are not recovered.

Transient errors during the computations to generate the time-frequency maps by X-SphRad prevented the full number of simulation to be produced. We note that the set of simulations added to the data at any given scale is not guaranteed to be exactly the same, which explains small fluctuations at minimum and maximum recovery rate. This also bias the % median recovery per scale, as the missing simulation is considered 'not recovered'. At distance ~ 100 Mpc, the fraction of recovered injections by the DoG method is consistently more sensitive than the current clustering implementation in X-SphRad with an increase by a factor of ~ 2 in efficiency, which potentially increase the detectable volume by nearly an order of magnitude. Moreover, the median amount of pixel recovered for each simulation is also doubled by the DoG. Recovering more pixels of candidate GW events make them more likely to pass the glitch rejection thresholds and cuts (See discussion on the correlation matrix in Chapter 3).

The levelling off of the recovery rate at the furthest distances is due to glitch background. As discussed in previous chapters, loud glitches are common and tend to overlap with simulated signals on the time-frequency map. The resulting candidate event is thus falsely treated as recovered. Difference between the baseline of the two models is explained by the longer duration of the ADI-D waveform, and the higher likelihood of a glitch happening at the same time.

The method is promising and shows great potential. Full assessment of the power of the method requires applying the same treatment to background candidate events. Efforts are underway to assess the impact of the different candidate event selection on the background rejection cuts & thresholds by implementing the method in X-SphRad.

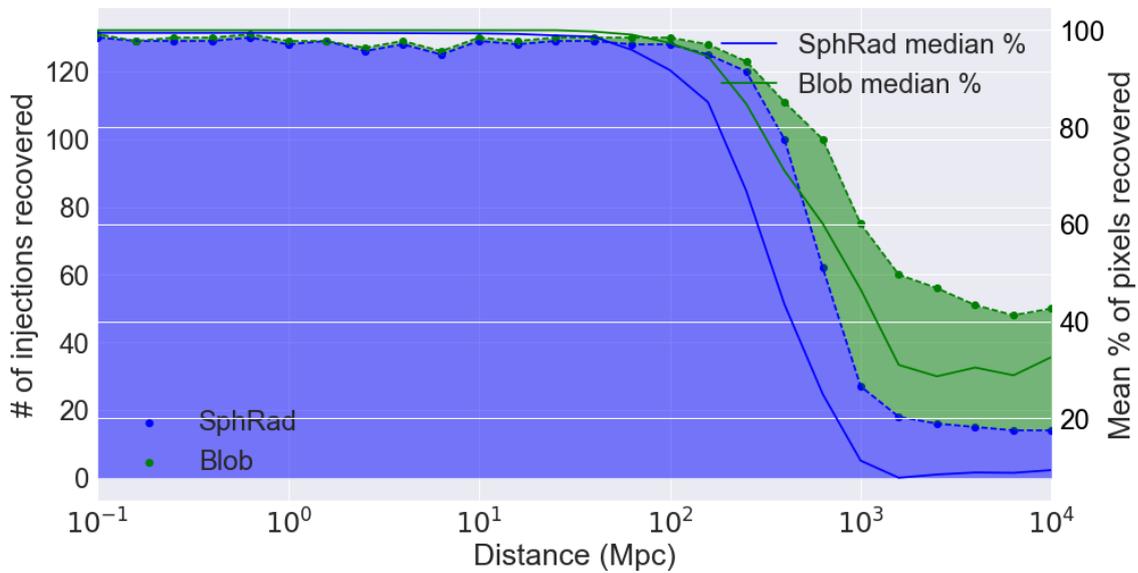


Figure 4.6: Comparison of the current X-SphRad clustering with the DoG for the ADI-E waveform. The DoG method was consistently more sensitive than the current clustering implementation in X-SphRad at distance above 500 Mpc. The DoG recovers about twice as many pixels per candidate event as X-SphRad at the same distance. More information as discussed in text.

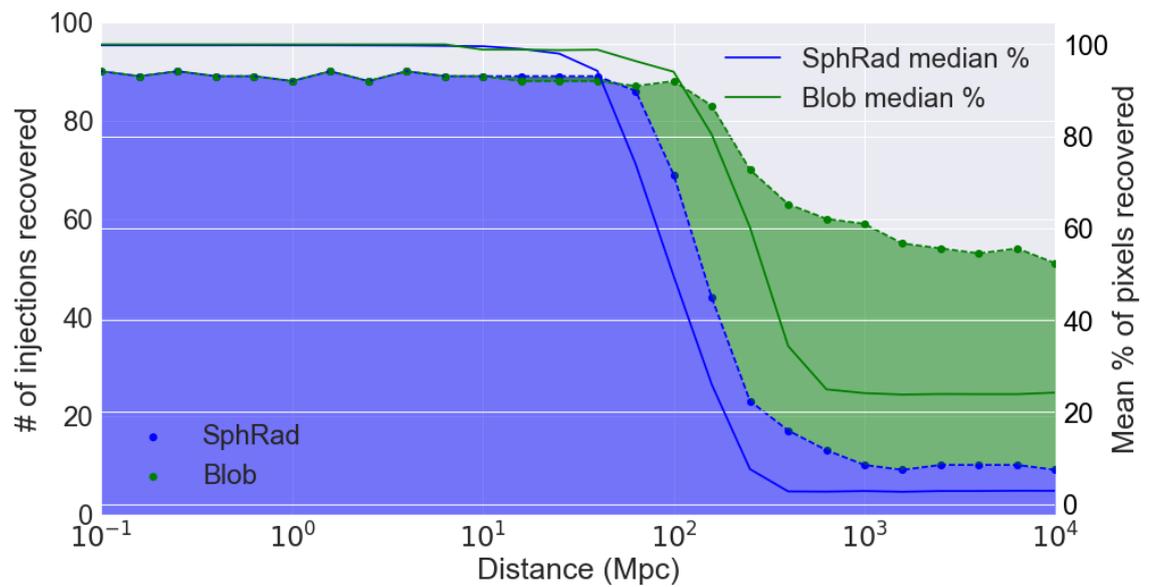


Figure 4.7: Comparison of the current X-SphRad clustering with the DoG for the ADI-D waveform. The DoG method was consistently more sensitive than the current clustering implementation in X-SphRad at distance above 100 Mpc. The DoG recovers about twice as many pixels per candidate event as X-SphRad at the same distance. More information as discussed in text.

CONCLUSION

Searches for long-duration GWs typically scan all the available data for signals arriving at any time from any direction on the sky. Because of the large parameter space covered, searches require model-independent processing and fast analysis techniques. For my PhD thesis, I integrated a set of fast cross-correlation routines in the spherical harmonic domain (SphRad) [50] into X-pipeline [95], a targeted GW search pipeline commonly used to search for GW counterparts of short and long duration GRBs & core-collapse supernovae.

Transient GW searches perform a coherent analysis on a linear combination of the detector output, weighted by the detector antenna response and noise spectrum. As the combination is chosen to maximise the Signal-to-Noise Ratio of a GW with given polarisation, it naturally depends on the source sky location. Current all-sky searches deal with sky location uncertainty by mapping a grid of points on the sky, and evaluating key statistics iteratively over the whole grid. The overall search speed thus depends approximately linearly on the number of points on the sky necessary to yield the desired accuracy.

Spherical harmonic decomposition used in this context allows for the sky position dependency of the coherent analysis to be isolated from the data [40] and cached for re-use, saving both time and processing units. Moreover, the spherical harmonic approach offers a fundamentally different view of the data, allowing for new possibilities for rejecting non-Gaussian background noise that could be mistaken for a GW signal. For example, the ratio of energy in the lowest harmonic mode ($l = 0$) to the

sum of the energies in the higher order modes ($l > 0$) is a powerful discriminator to discard instrumental noise transients.

The combined search pipeline, X-SphRad, underwent a thorough internal review within the LIGO collaboration, which I led. The pipeline good functioning was assessed by rigorous tests including comparing a test data set with a standard sky grid-based analysis, adding and recovering signals with calibration-line-like properties, and thorough follow-up investigations of noise artefacts and potential GW candidates. X-SphRad analysed the data taken during the first advanced LIGO observing run (O1) from September 2015 to January 2016. An observational results paper is in the final state of review by the LIGO-Virgo Collaboration [22].

By carefully actuating on the detector test masses, it is possible to add simulated GW signals to the LIGO data strain which allows us to characterise calibration uncertainties, unbiased analysis sensitivities and online-search responses with collaboration-wide impact. This is a key step in validation of candidate GW signals [6].

My current research focus is on innovation in time-frequency analysis. Current events-of-interest in searches for unmodelled GW bursts are flagged by selecting clusters of pixels on high-resolution spectrograms. Due to the presence of non-stationary non-Gaussian noise in the detectors, strong thresholding on the intensity is used to select seed-pixels from which clusters will grow. The thresholding method is efficient at discarding instrumental artefacts but requires that at least few pixels of a putative GW signal are louder than the set threshold, which can significantly limit the sensitivity of the search in non-stationary data.

We searched the data from the first advanced LIGO observing run for long-duration gravitational wave transients. We report a set of candidate events that passed the coherence and consistency tests. All candidates are consistent with the background and we report no long-duration gravitational wave detection. We thus set the 90% confidence limit on 6 astrophysically motivated signal templates.

I have developed a novel pixel clustering method that does not depend on the amplitude of potential signals. By using an edge detection algorithm, I quantify each pixel in the spectrogram by its similarity with its neighbours then extract features of sharply changing intensity (or ‘edge’). By ranking clusters of pixels by their ‘edginess’ instead of energy, the method assumes only that a potential signal need to be louder than the noise at a particular time/frequency bin and its close neighbours. The method has shown promising results in preliminary tests with a potential increase in observable volume by about a factor of 10. A simplified version of the algorithm was implemented in X-SphRad and large-scale testings are currently being processed.

Future work on X-SphRad will focus on reducing the latency of the analysis with end goal to allow for long-duration GW event reporting within the order of minutes.

CHAPTER 

APPENDIX A

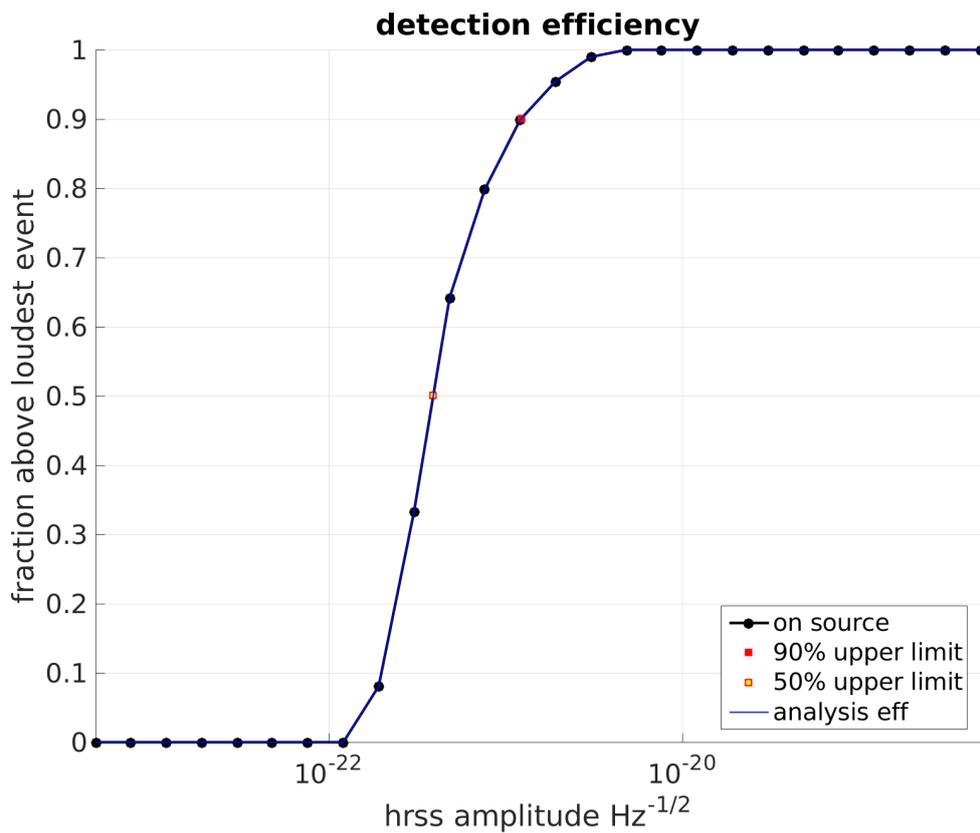


Figure 6.1: ADI-A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

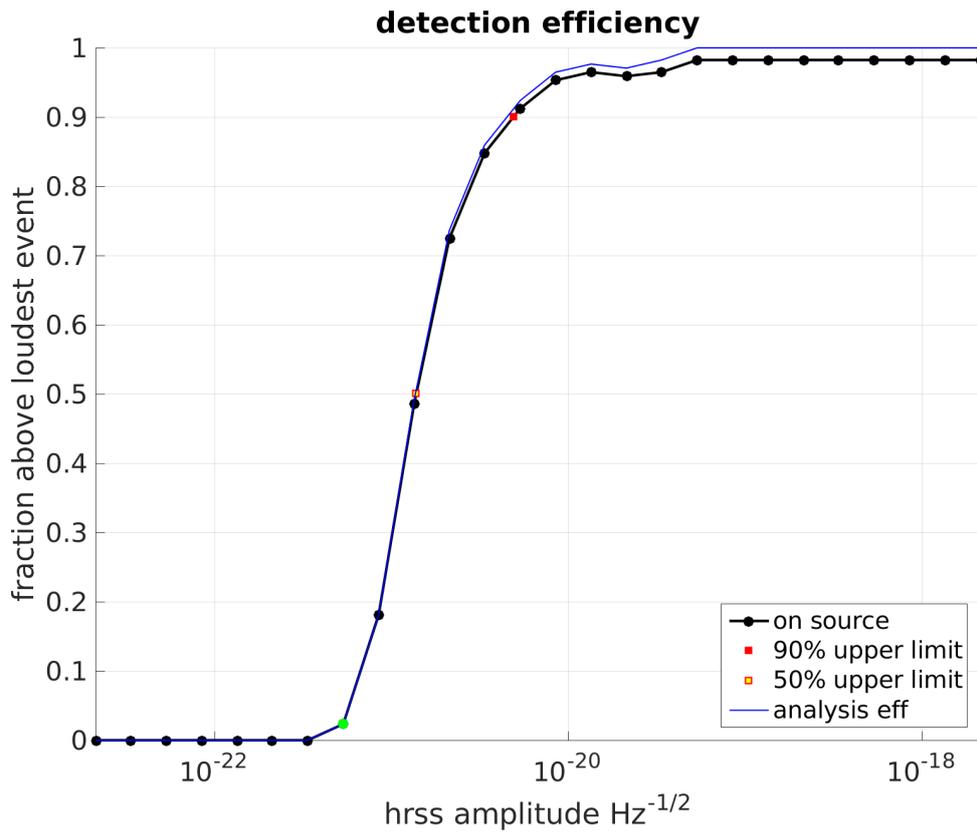


Figure 6.3: ADI-C: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

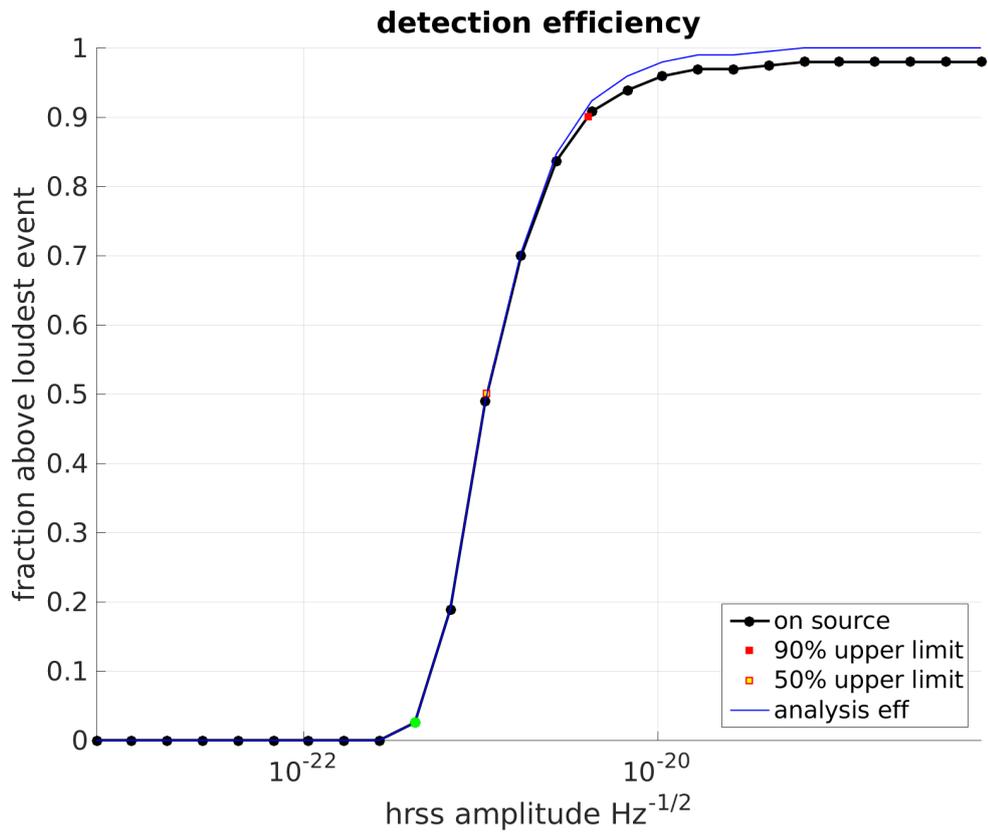


Figure 6.4: ADI-D: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

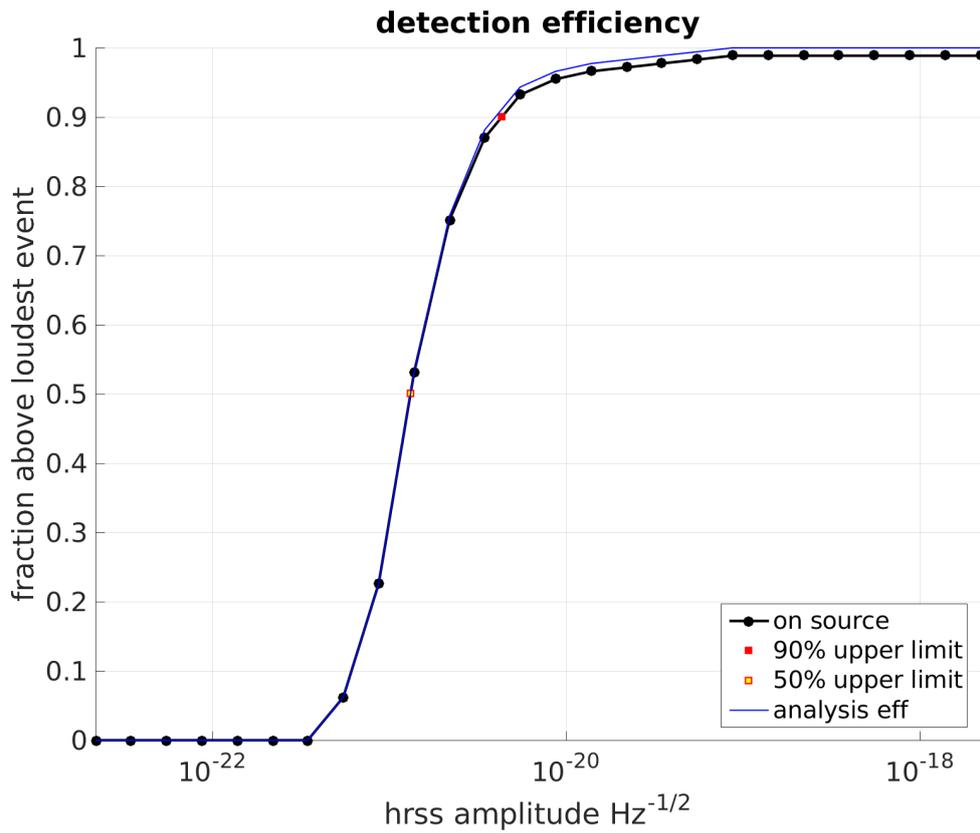


Figure 6.5: ADI-E: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

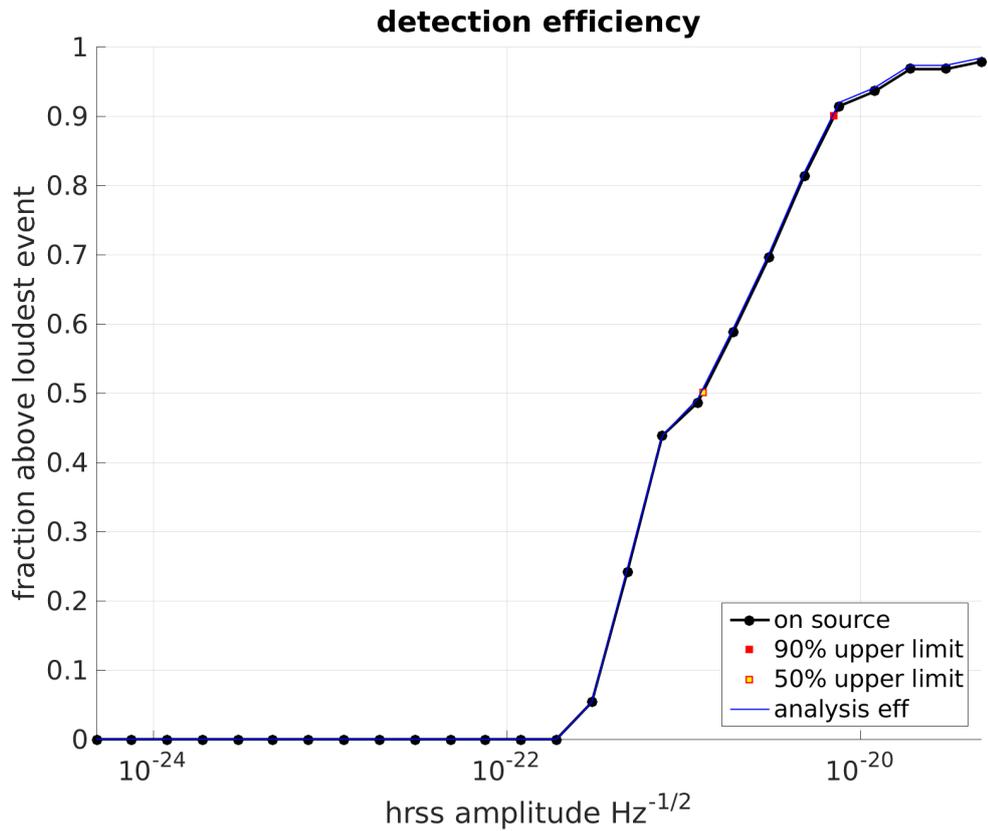


Figure 6.6: Magnetar F: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

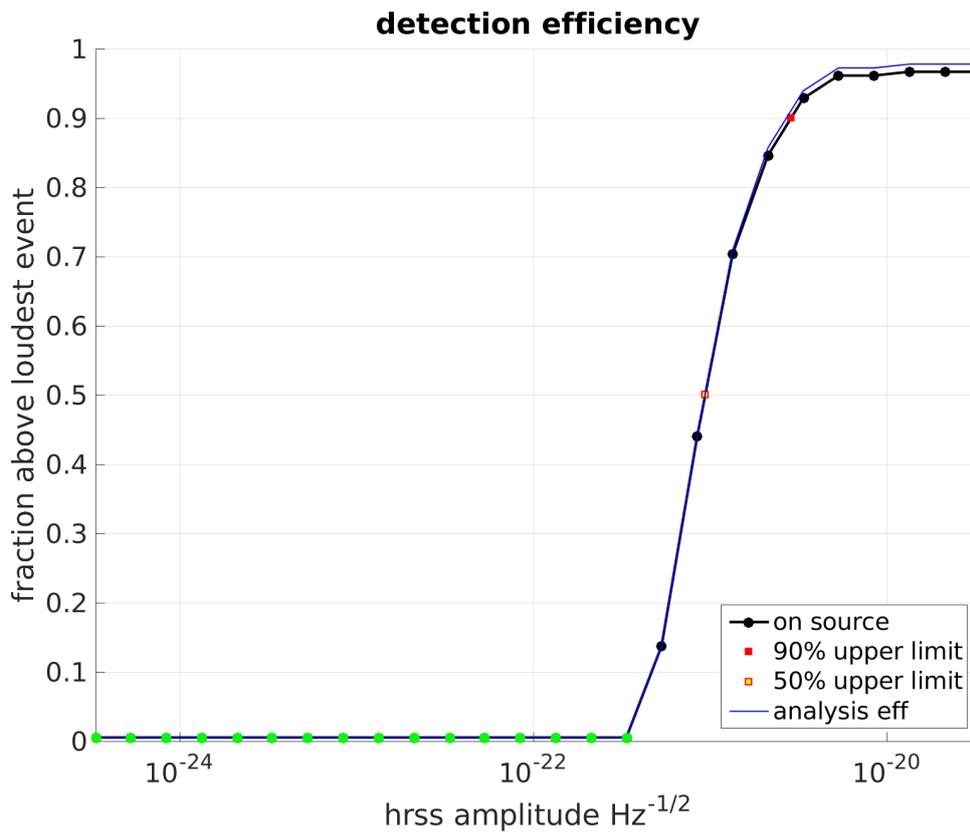


Figure 6.7: Magnetar G: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

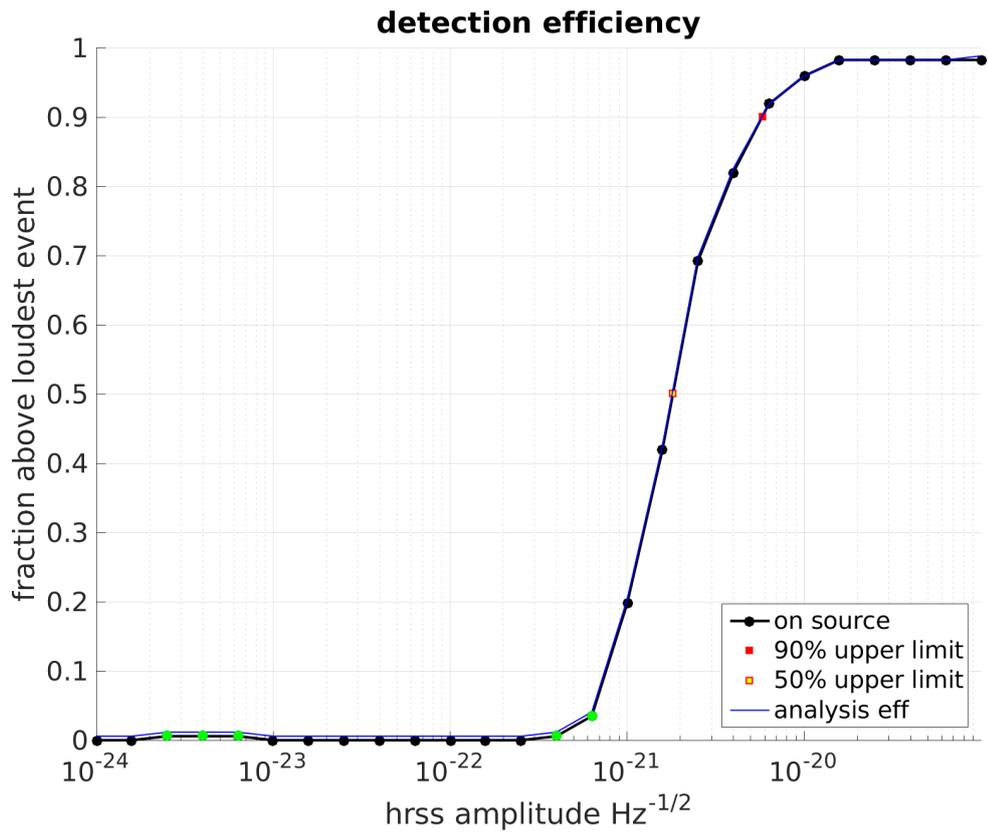


Figure 6.8: Quad-A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

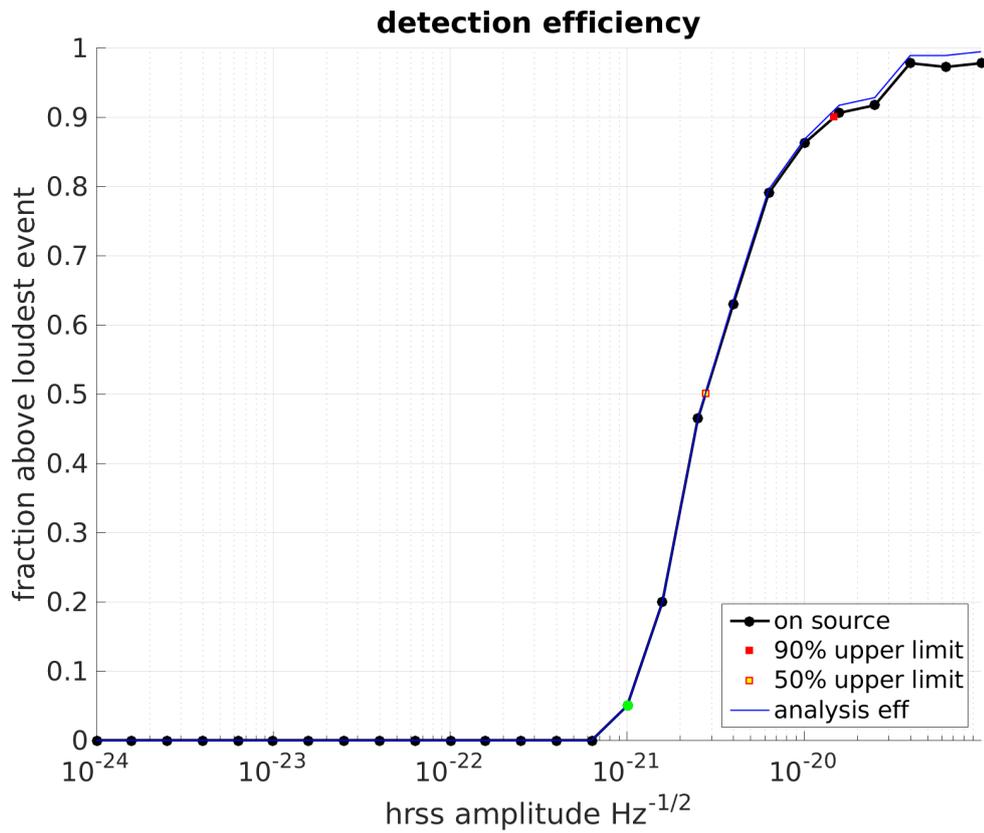


Figure 6.10: SG A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

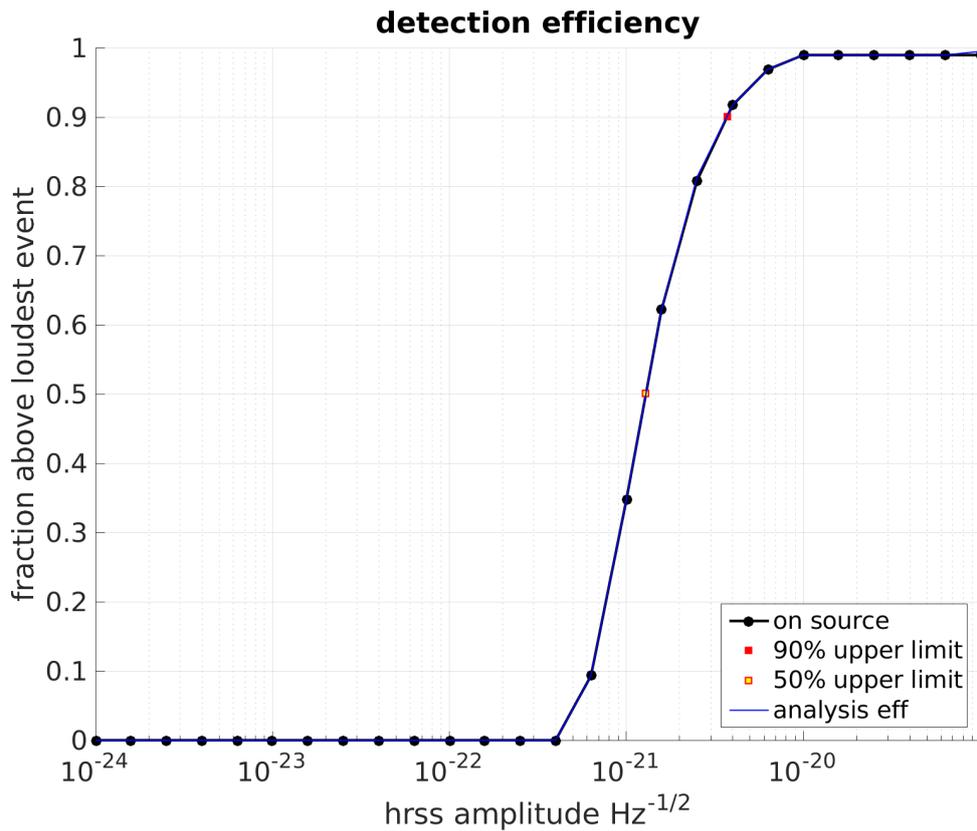


Figure 6.11: SG C: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled valued with $0 < \text{efficiency} < 5\%$. Blue curve show the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

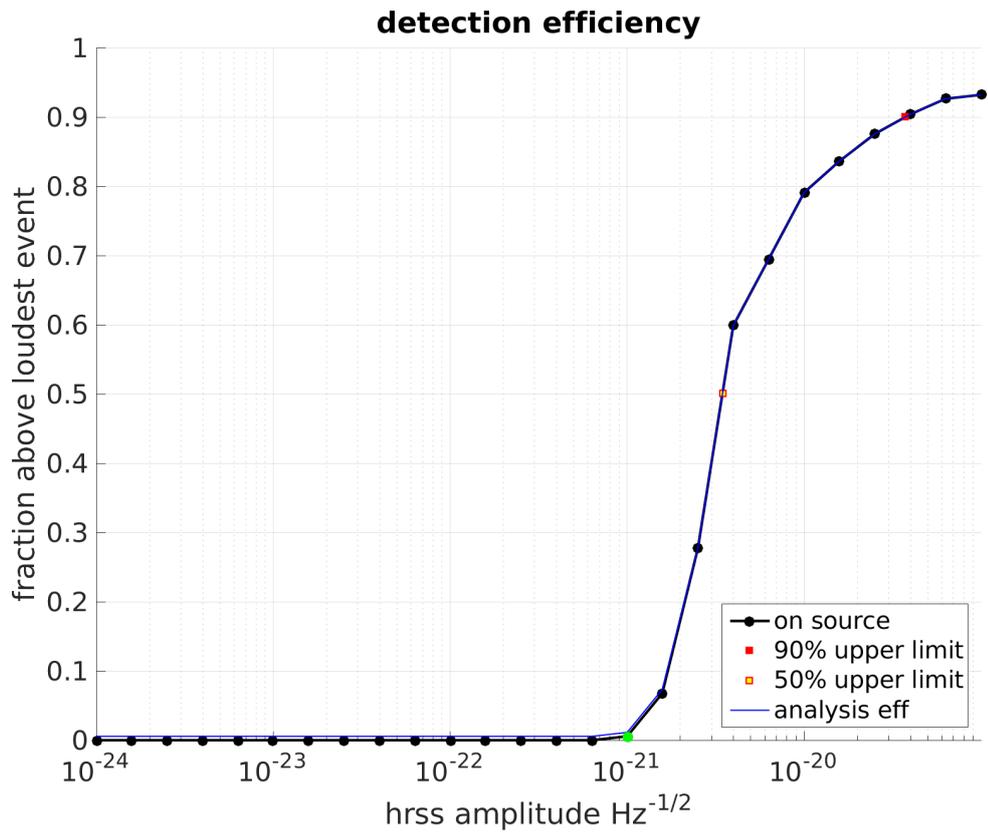


Figure 6.12: WNB A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

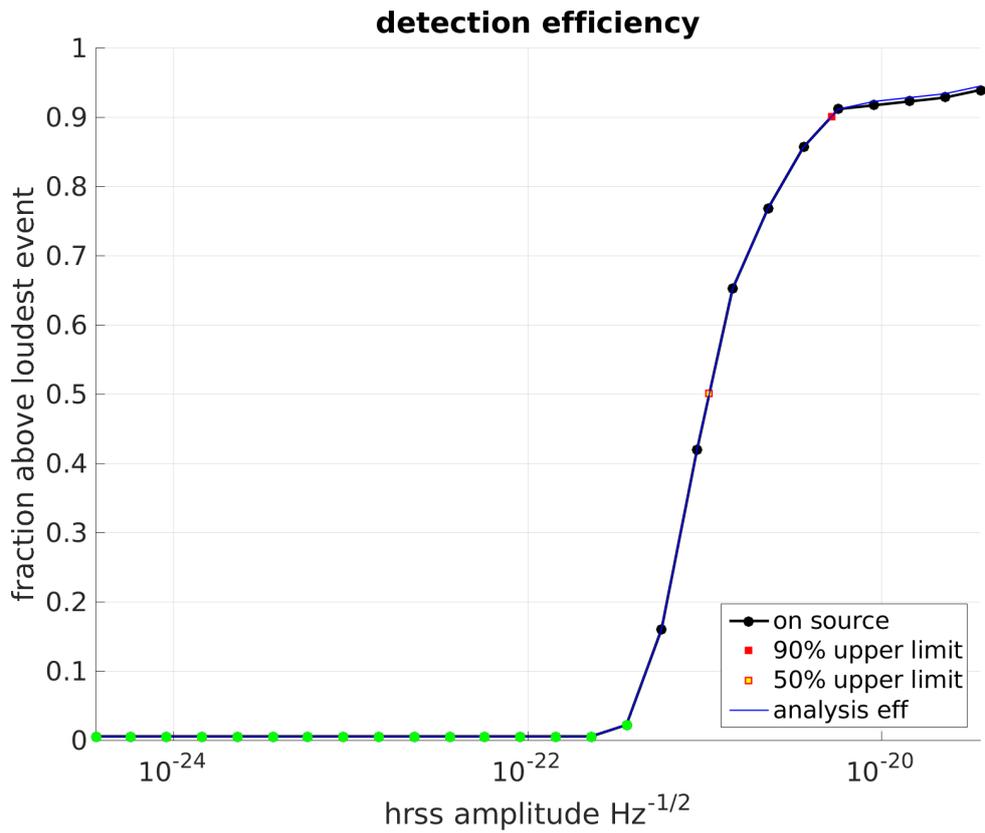


Figure 6.13: WNB B: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

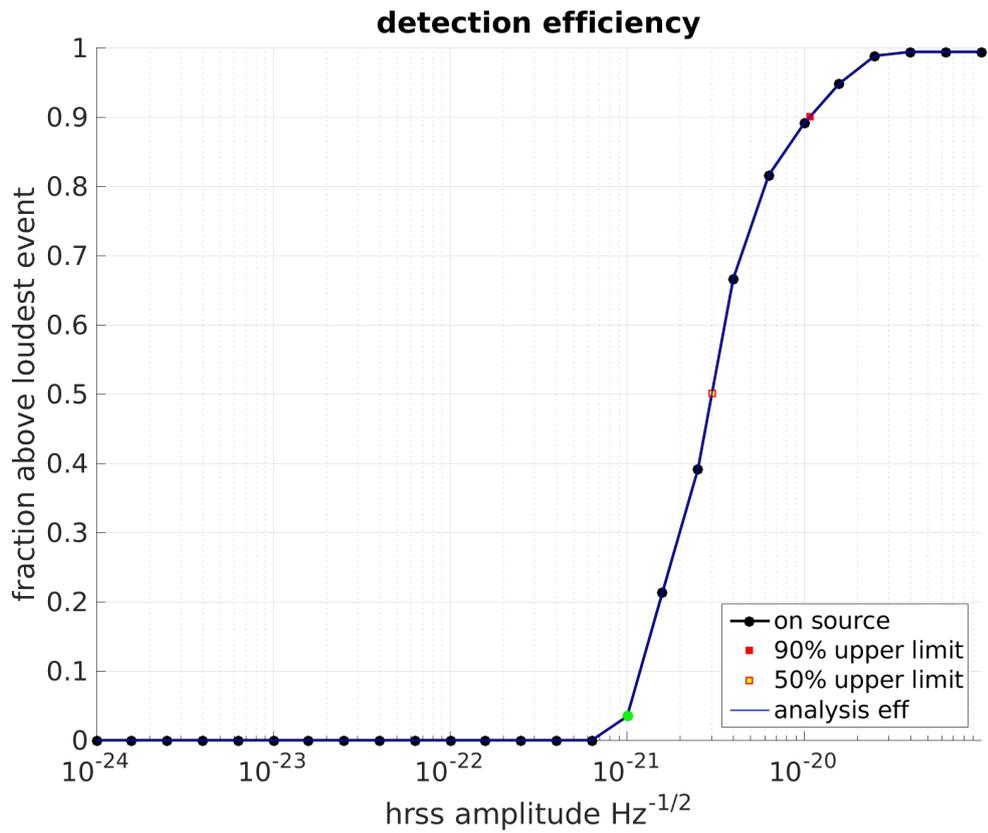


Figure 6.14: LINE A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

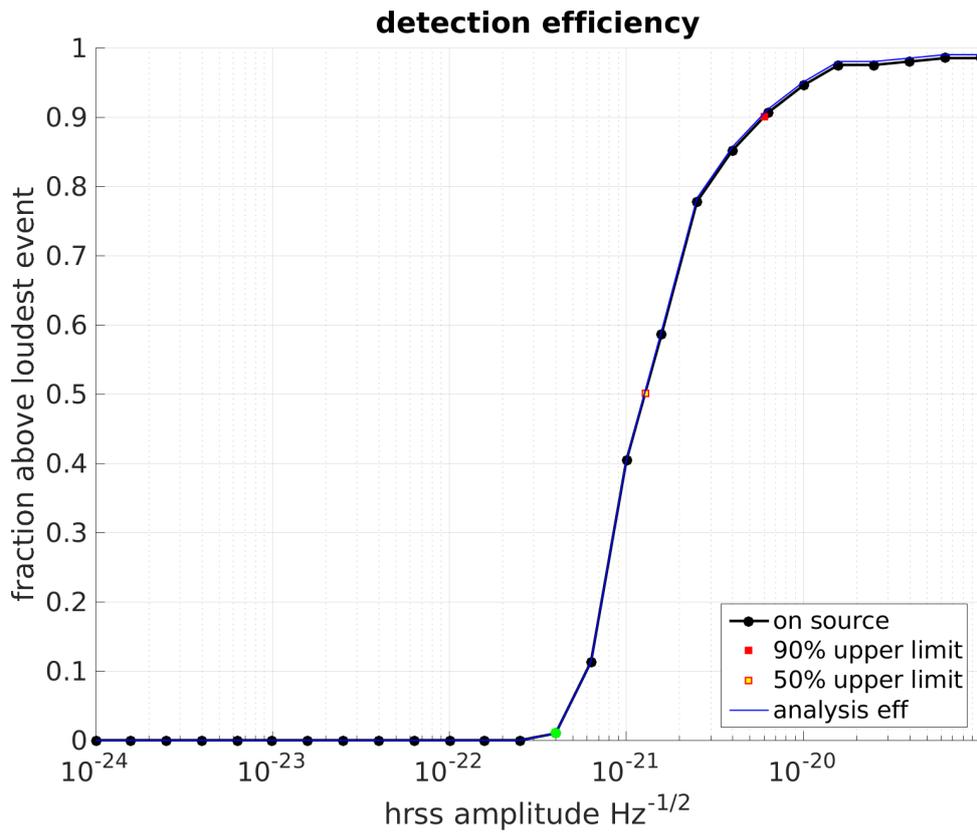


Figure 6.15: Line B: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

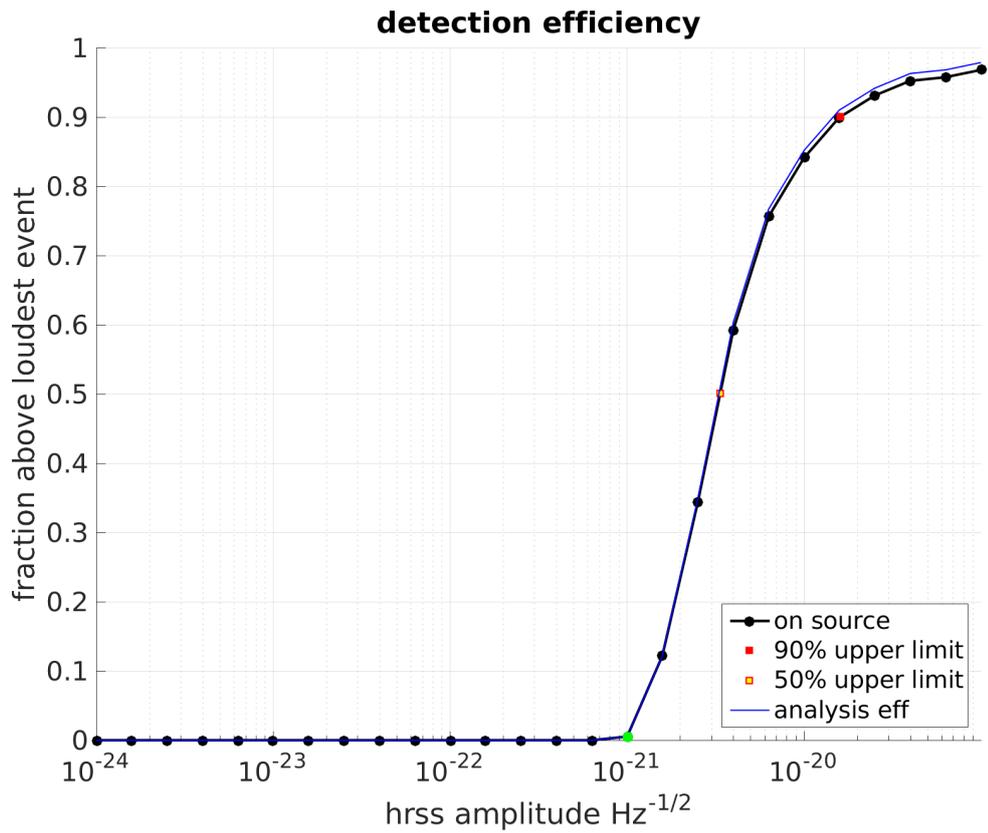


Figure 6.16: Mono A: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

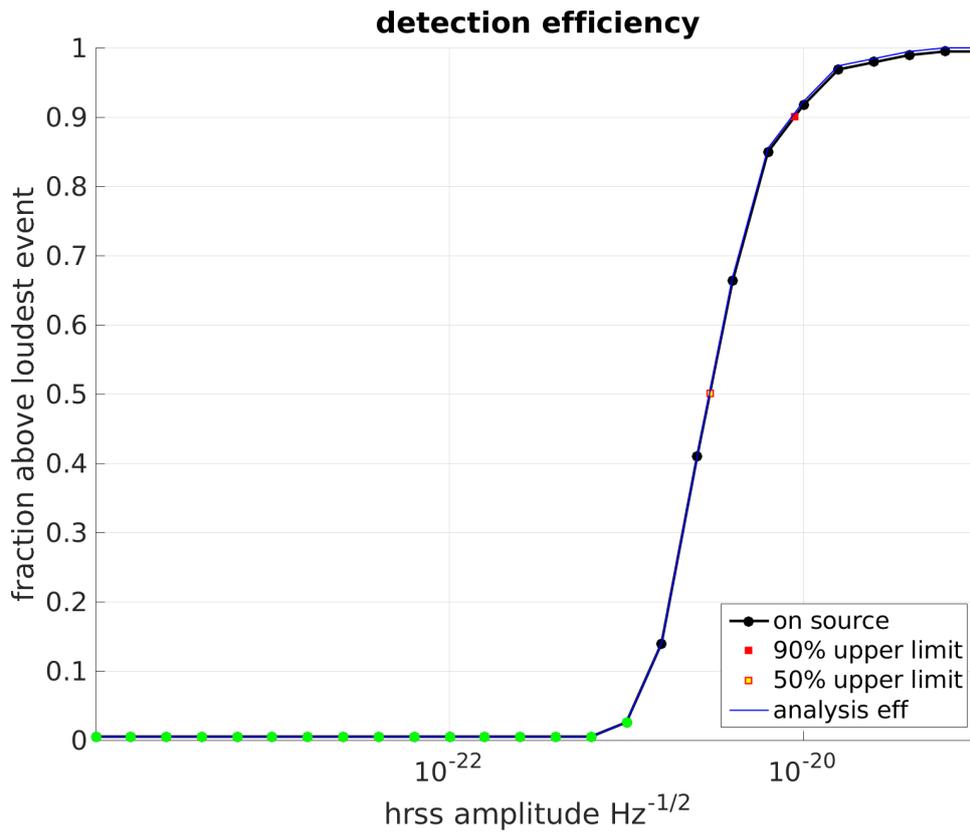


Figure 6.17: Mono C: Fraction of injections recovered with significance greater than loudest event in (dummy) on-source. Black dots are sampled values, red and yellow dot is respectively the 90% and 50% efficiency obtained from interpolation. Green dots mark sampled values with $0 < \text{efficiency} < 5\%$. Blue curve shows the efficiency when DQ flags are not applied to injections. We obtain the 90% upper limit from fit

BIBLIOGRAPHY

- [1] Aasi, J., Abadie, J., Abbott, B. P., et al. Search for long-lived gravitational-wave transients coincident with long gamma-ray bursts. *Physical Review D*, 88(12):122004, dec 2013. ISSN 1550-7998. doi: 10.1103/PhysRevD.88.122004.
- [2] Aasi, J., Abbott, B. P., Abbott, R., et al. Search for Gravitational Waves Associated with gamma-ray Bursts Detected by the Interplanetary Network. *Physical Review Letters*, 113(1):011102, jun 2014. ISSN 0031-9007. doi: 10.1103/PhysRevLett.113.011102.
- [3] Aasi, J., Abbott, B. P., Abbott, R., et al. Search for Gravitational Waves Associated with Gamma-ray Bursts Detected by the Interplanetary Network. *Physical Review Letters*, 113(1):011102, jun 2014. ISSN 0031-9007. doi: 10.1103/PhysRevLett.113.011102.
- [4] Aasi, J., Abbott, B. P., Abbott, R., et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, apr 2015. ISSN 0264-9381. doi: 10.1088/0264-9381/32/7/074001.
- [5] Aasi, J., Abbott, B. P., Abbott, R., et al. First low frequency all-sky search for continuous gravitational wave signals. *Physical Review D*, 93(4):042007, feb 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.93.042007.

BIBLIOGRAPHY

- [6] Aasi, J., Abbott, B. P., Abbott, R., et al. The LSC–Virgo White Paper on Gravitational Wave Searches and Astrophysics (2014–2015 edition). *Living Reviews in Relativity*, 19:1–39, 2016. ISSN 14338351. doi: 10.1007/lrr-2016-1.
- [7] Abadie, J., Abbott, B. P., Abbott, R., et al. Search for gravitational waves from compact binary coalescence in LIGO and Virgo data from S5 and VSR1. *Physical Review D*, 82(10):102001, nov 2010. ISSN 1550-7998. doi: 10.1103/PhysRevD.82.102001.
- [8] Abadie, J., Abbott, B. P., Abbott, R., et al. Directional Limits on Persistent Gravitational Waves Using LIGO S5 Science Data. *Physical Review Letters*, 107(27):271102, dec 2011. ISSN 0031-9007. doi: 10.1103/PhysRevLett.107.271102.
- [9] Abbott, B., Abbott, R., Adhikari, R., et al. First upper limits from LIGO on gravitational wave bursts. *Physical Review D*, 69(10):102001, may 2004. ISSN 1550-7998. doi: 10.1103/PhysRevD.69.102001.
- [10] Abbott, B., Abbott, R., Adhikari, R., et al. Searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run. *Physical Review D*, 76(8):082001, oct 2007. ISSN 1550-7998. doi: 10.1103/PhysRevD.76.082001.
- [11] Abbott, B., Abbott, R., Adhikari, R. X., et al. Search for gravitational-wave bursts in LIGO data from the fourth science run. *Classical and Quantum Gravity*, 24(22):5343–5369, 2007.
- [12] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914.

- Classical and Quantum Gravity*, 33(13):134001, jul 2016. ISSN 0264-9381. doi: 10.1088/0264-9381/33/13/134001.
- [13] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116(6): 061102, feb 2016. doi: 10.1103/PhysRevLett.116.061102.
- [14] Abbott, B. P., Abbott, R., Abbott, T. D., et al. GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence. *Physical Review Letters*, 116(24):241103, jun 2016. ISSN 0031-9007. doi: 10.1103/PhysRevLett.116.241103.
- [15] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Upper limits on the rates of binary neutron star and neutron-star–black-hole mergers from Advanced LIGO’s first observing run. *The Astrophysical Journal*, 832(2):L21, nov 2016. ISSN 2041-8213. doi: 10.3847/2041-8205/832/2/L21.
- [16] Abbott, B. P., Abbott, R., Abbott, T. D., et al. First targeted search for gravitational-wave bursts from core-collapse supernovae in data of first-generation laser interferometer detectors. *Physical Review D*, 94(10):102001, nov 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.94.102001.
- [17] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Search for transient gravitational waves in coincidence with short-duration radio transients during 2007–2013. *Physical Review D*, 93(12):122008, jun 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.93.122008.

BIBLIOGRAPHY

- [18] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Binary Black Hole Mergers in the First Advanced LIGO Observing Run. *Physical Review X*, 6(4):041015, oct 2016. ISSN 2160-3308. doi: 10.1103/PhysRevX.6.041015.
- [19] Abbott, B. P., Abbott, R., Abbott, T. D., et al. GW150914: First results from the search for binary black hole coalescence with Advanced LIGO. *Physical Review D*, 93(12):122003, jun 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.93.122003.
- [20] Abbott, B. P., Abbott, R., Abernathy, M. R., et al. Observing gravitational-wave transient GW150914 with minimal assumptions. *Physical Review D*, 93(12): Art. No. 122004, 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.93.122004.
- [21] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Search for Gravitational Waves Associated with Gamma-Ray Bursts during the First Advanced LIGO Observing Run and Implications for the Origin of GRB 150906B. *The Astrophysical Journal*, 841(2):89, may 2017. ISSN 1538-4357. doi: 10.3847/1538-4357/aa6c47.
- [22] Abbott, B. P., Abbott, R., Abbott, T. D., et al. All-sky search for short gravitational-wave bursts in the first Advanced LIGO run. *Physical Review D*, 95(4):042003, feb 2017. ISSN 2470-0010. doi: 10.1103/PhysRevD.95.042003.
- [23] Abbott, B. P., Abbott, R., Abbott, T. D., et al. All-sky search for short gravitational-wave bursts in the first Advanced LIGO run. *Physical Review D*, 95(4):042003, feb 2017. ISSN 2470-0010. doi: 10.1103/PhysRevD.95.042003.
- [24] Abbott, B. P., Abbott, R., Abbott, T. D., et al. The LSC-Virgo White Paper on Gravitational Wave Data Analysis and Astrophysics (July 2017 edition). Technical report, 2017.

- [25] Abbott, B. P., Abbott, R., Abbott, T. D., et al. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. *Physical Review D*, 95(6):062003, mar 2017. ISSN 2470-0010. doi: 10.1103/PhysRevD.95.062003.
- [26] Abbott, B. P., Abbott, R., Abbott, T. D., et al. GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2. *Physical Review Letters*, 118(22):221101, jun 2017. ISSN 0031-9007. doi: 10.1103/PhysRevLett.118.221101.
- [27] Ali, A., Ahmad, S., Nawaz, M., Ullah, S., and Aqeel, M. Bayesian Inference on Gravitational Waves. *Pakistan Journal of Statistics and Operation Research*, 11(4):645, dec 2015. ISSN 2220-5810. doi: 10.18187/pjsor.v11i4.1053.
- [28] Ali, M. and Clausi, D. Using the Canny edge detector for feature extraction and enhancement of remote sensing images. In *IGARSS 2001. Scanning the Present and Resolving the Future. Proceedings. IEEE 2001 International Geoscience and Remote Sensing Symposium (Cat. No.01CH37217)*, volume 5, pages 2298–2300. IEEE. ISBN 0-7803-7031-7. doi: 10.1109/IGARSS.2001.977981.
- [29] Anderson, W. G., Brady, P. R., Creighton, J. D. E., and Flanagan, E. E. A power filter for the detection of burst sources of gravitational radiation in interferometric detectors. *International Journal of Modern Physics D*, 09(03): 303–307, jan 2000. ISSN 0218-2718. doi: 10.1142/S0218271800000323.
- [30] Anderson, W. G., Brady, P. R., Creighton, J. D. E., and Flanagan, É. É. Excess power statistic for detection of burst sources of gravitational radi-

BIBLIOGRAPHY

- ation. *Physical Review D*, 63(4):042003, jan 2001. ISSN 0556-2821. doi: 10.1103/PhysRevD.63.042003.
- [31] Arabie, P., Hubert, L. J., and De Soete, G. *Clustering and Classification*. WORLD SCIENTIFIC, jan 1996. ISBN 978-981-02-1287-2. doi: 10.1142/1930.
- [32] Arai, K. LIGO-G1501380-v1: Ideas for RF Glitch Identification / Mitigation, 2015.
- [33] Assirati, L., Silva, N. R., Berton, L., Lopes, A. A., and Bruno, O. M. Performing edge detection by Difference of Gaussians using q-Gaussian kernels. *Journal of Physics: Conference Series*, 490(1):012020, mar 2014. ISSN 1742-6596. doi: 10.1088/1742-6596/490/1/012020.
- [34] Basu, M. Gaussian-based edge-detection methods-a survey. *IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews)*, 32(3):252–260, aug 2002. ISSN 1094-6977. doi: 10.1109/TSMCC.2002.804448.
- [35] Biscans, S., Warner, J., Mittleman, R., et al. Control strategy to limit duty cycle impact of earthquakes on the LIGO gravitational-wave detectors. *Classical and Quantum Gravity*, 35(5):055004, mar 2018. ISSN 0264-9381. doi: 10.1088/1361-6382/aaa4aa.
- [36] Bourne, R. *Fundamentals of digital imaging in medicine*. Springer, 2010. ISBN 1848820879.
- [37] B.P., F., Teukolski, S., and W.T., V. Numerical Recipes: the Art of Scientific Computing. pages 381–449. Cambridge University Press, 1986. ISBN 9780521880688.

-
- [38] Brady, P. R., Creighton, J. D. E., and Wiseman, A. G. Upper limits on gravitational-wave signals based on loudest events. *Classical and Quantum Gravity*, 21(20):S1775–S1781, oct 2004. ISSN 0264-9381. doi: 10.1088/0264-9381/21/20/020.
- [39] Cadonati, L. Coherent waveform consistency test for LIGO burst candidates. *Classical and Quantum Gravity*, 21(20):S1695–S1703, oct 2004. ISSN 0264-9381. doi: 10.1088/0264-9381/21/20/012.
- [40] Cannon, K. C. and C., K. Efficient algorithm for computing the time-resolved full-sky cross power in an interferometer with omnidirectional elements. *Physical Review D*, 75(12):123003, jun 2007. ISSN 1550-7998. doi: 10.1103/PhysRevD.75.123003.
- [41] Canny, J. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, nov 1986. ISSN 0162-8828. doi: 10.1109/TPAMI.1986.4767851.
- [42] Chassande-Mottin, E., Lebigot, E., Magaldi, H., et al. Wavelet graphs for the direct detection of gravitational waves. oct 2017.
- [43] Chatterji, S., Blackburn, L., Martin, G., and Katsavounidis, E. Multiresolution techniques for the detection of gravitational-wave bursts. *Classical and Quantum Gravity*, 21(20):S1809–S1818, oct 2004. ISSN 0264-9381. doi: 10.1088/0264-9381/21/20/024.
- [44] Christensen, N. LIGO S6 detector characterization studies. *Classical and Quantum Gravity*, 27(19):194010, oct 2010. ISSN 0264-9381. doi: 10.1088/0264-9381/27/19/194010.

BIBLIOGRAPHY

- [45] Coughlin, M., Meyers, P., Thrane, E., Luo, J., and Christensen, N. Detectability of eccentric compact binary coalescences with advanced gravitational-wave detectors. *Physical Review D*, 91(6):063004, mar 2015. ISSN 1550-7998. doi: 10.1103/PhysRevD.91.063004.
- [46] Coughlin, M. W. Identification of long-duration noise transients in LIGO and Virgo. *Classical and Quantum Gravity*, 28(23):235008, dec 2011. ISSN 0264-9381. doi: 10.1088/0264-9381/28/23/235008.
- [47] Deriche, R. Using Canny's Criteria to Derive a Recursively Implemented Optimal Edge Detector. *International Journal of Computer Vision*, pages 167–187, 1987.
- [48] Dove, E. L., Philip, K., Gotteiner, N. L., et al. A method for automatic edge detection and volume computation of the left ventricle from ultrafast computed tomographic images. *Investigative radiology*, 29(11):945–54, nov 1994. ISSN 0020-9996.
- [49] Duda, R. O. and Hart, P. E. P. E. *Pattern classification and scene analysis*. Wiley, 1973. ISBN 0471223611.
- [50] Edwards, M. *On the search for intermediate duration gravitational waves using the spherical harmonic basis*. PhD thesis, 2013.
- [51] Edwards, M. and Sutton, P. J. A new glitch-rejection algorithm forged in the spherical harmonic basis. *Journal of Physics: Conference Series*, 363(1): 012025, jun 2012. ISSN 1742-6596. doi: 10.1088/1742-6596/363/1/012025.

- [52] Einstein, A. *Relativity : the special and the general theory; a popular exposition*. Martino Publishing, 2010. ISBN 9781891396304.
- [53] Feng-ying Cui, Li-jun Zou, and Bei Song. Edge feature extraction based on digital image processing techniques. In *2008 IEEE International Conference on Automation and Logistics*, pages 2320–2324. IEEE, sep 2008. ISBN 978-1-4244-2502-0. doi: 10.1109/ICAL.2008.4636554.
- [54] Finn, L. S. and Chernoff, D. F. Observing binary inspiral in gravitational radiation: One interferometer. *Physical Review D*, 47(6):2198–2219, mar 1993. doi: 10.1103/PhysRevD.47.2198.
- [55] Flanagan, É. É. and Hughes, S. A. Measuring gravitational waves from binary black hole coalescences. I. Signal to noise for inspiral, merger, and ringdown. *Physical Review D*, 57(8):4535–4565, apr 1998. ISSN 0556-2821. doi: 10.1103/PhysRevD.57.4535.
- [56] Franco, S. *Searching for long transient gravitational waves in the LIGO-Virgo data*. PhD thesis, jul 2014.
- [57] Giacomazzo, B. and Perna, R. Formation of stable magnetars from binary neutron star mergers. *The Astrophysical Journal*, 771(2):L26, jun 2013. ISSN 2041-8205. doi: 10.1088/2041-8205/771/2/L26.
- [58] Gossan, S. E., Sutton, P., Stuver, A., et al. Observing gravitational waves from core-collapse supernovae in the advanced detector era. *Physical Review D*, 93(4):042002, feb 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.93.042002.

BIBLIOGRAPHY

- [59] Gráfová, L., Kasparová, M., Kakawand, S., Procházka, A., and Dostálová, T. Study of edge detection task in dental panoramic radiographs. *Dento maxillo facial radiology*, 42(7):20120391, 2013. ISSN 0250-832X. doi: 10.1259/dmfr.20120391.
- [60] Gürsel, Y. and Tinto, M. Near optimal solution to the inverse problem for gravitational-wave bursts. *Physical Review D*, 40(12):3884–3938, dec 1989. ISSN 0556-2821. doi: 10.1103/PhysRevD.40.3884.
- [61] Hartle, J. B. J. B. *Gravity : an introduction to Einstein's general relativity*. Addison-Wesley, 2003. ISBN 0805386629.
- [62] Helstrom, C. W. *Statistical theory of signal detection*. Pergamon Press, 1968. ISBN 0080132650.
- [63] Honda, R., Yamagishi, S., and Kanda, N. Astrophysically motivated time–frequency clustering for burst gravitational wave search: application to TAMA300 data. *Classical and Quantum Gravity*, 25(18):184035, sep 2008. ISSN 0264-9381. doi: 10.1088/0264-9381/25/18/184035.
- [64] Horn, B. *Robot vision*. MIT Press, 1986. ISBN 9780262081597.
- [65] Hulse, R. A. and Taylor, J. H. Discovery of a pulsar in a binary system. *The Astrophysical Journal*, 195:L51, jan 1975. ISSN 0004-637X. doi: 10.1086/181708.
- [66] Husa, S., Khan, S., Hannam, M., et al. Frequency-domain gravitational waves from nonprecessing black-hole binaries. I. New numerical waveforms and

- anatomy of the signal. *Physical Review D*, 93(4):044006, feb 2016. ISSN 2470-0010. doi: 10.1103/PhysRevD.93.044006.
- [67] Isogai, T., Collaboration, t. L. S. C. a., and Collaboration, V. Used percentage veto for LIGO and virgo binary inspiral searches. *Journal of Physics: Conference Series*, 243(1):012005, aug 2010. ISSN 1742-6596. doi: 10.1088/1742-6596/243/1/012005.
- [68] Jassim, F. A. Semi-Optimal Edge Detector based on Simple Standard Deviation with Adjusted Thresholding. *International Journal of Computer Applications*, 68(2):975–8887, 2013.
- [69] Jhurani, C. and Demkowicz, L. Multiscale modeling using goal-oriented adaptivity and numerical homogenization. Part I: Mathematical formulation and numerical results. *Computer Methods in Applied Mechanics and Engineering*, 213-216:399–417, mar 2012. ISSN 0045-7825. doi: 10.1016/J.CMA.2011.06.011.
- [70] Karki, S., Tuyenbayev, D., Kandhasamy, S., et al. The Advanced LIGO photon calibrators. *Review of Scientific Instruments*, 87(11):114503, nov 2016. ISSN 0034-6748. doi: 10.1063/1.4967303.
- [71] Klimenko, S., Yakushin, I., Mercer, A., and Mitselmakher, G. A coherent method for detection of gravitational wave bursts. *Classical and Quantum Gravity*, 25(11):114029, jun 2008. ISSN 0264-9381. doi: 10.1088/0264-9381/25/11/114029.
- [72] Kokkotas, K. D. Gravitational wave physics. *Encyclopedia of Physical Science and Technology*, 7, 2002.

BIBLIOGRAPHY

- [73] Lakhani, K., Minocha, B., and Gugnani, N. Analyzing edge detection techniques for feature extraction in dental radiographs. *Perspectives in Science*, 8: 395–398, sep 2016. ISSN 22130209. doi: 10.1016/j.pisc.2016.04.087.
- [74] Loveridge, L. C. Physical and Geometric Interpretations of the Riemann Tensor, Ricci Tensor, and Scalar Curvature. jan 2004.
- [75] Marr, D. and Hildreth, E. Theory of edge detection. *Proceedings of the Royal Society of London. Series B, Biological sciences*, 207(1167):187–217, feb 1980. ISSN 0950-1193. doi: 10.1098/RSPB.1980.0020.
- [76] Mcleod, D. GWpy: a python package for gravitational-wave astrophysics., 2017.
- [77] Misner, C. W., Thorne, K. S., and Wheeler, J. A. *Gravitation*. ISBN 0716703440.
- [78] Moon, H., Chellappa, R., and Rosenfeld, A. Optimal edge-based shape detection. *IEEE Transactions on Image Processing*, 11(11):1209–1227, nov 2002. ISSN 1057-7149. doi: 10.1109/TIP.2002.800896.
- [79] Oppenheim, A. V. and Schaffer, R. W. *Discrete-time signal processing*. Pearson, 2010. ISBN 0131988425.
- [80] Ott, C., OConnor, E., Gossan, S., et al. Core-Collapse Supernovae, Neutrinos, and Gravitational Waves. *Nuclear Physics B - Proceedings Supplements*, 235-236:381–387, feb 2013. ISSN 09205632. doi: 10.1016/j.nuclphysbps.2013.04.036.

- [81] Ott, C. D. The gravitational-wave signature of core-collapse supernovae. *Classical and Quantum Gravity*, 26(6):063001, mar 2009. ISSN 0264-9381. doi: 10.1088/0264-9381/26/6/063001.
- [82] Öztürk, and Akdemir, B. Comparison of Edge Detection Algorithms for Texture Analysis on Glass Production. *Procedia - Social and Behavioral Sciences*, 195:2675–2682, jul 2015. ISSN 1877-0428. doi: 10.1016/J.SBSPRO.2015.06.477.
- [83] Palamos, J. aLIGO LHO Logbook. Technical report, 2015.
- [84] Pescalli, A., Ghirlanda, G., Salvaterra, R., et al. The rate and luminosity function of long Gamma Ray Bursts. *Astronomy & Astrophysics, Volume 587, id.A40, 11 pp.*, 587, jun 2015. doi: 10.1051/0004-6361/201526760.
- [85] Piro, A. L. and Thrane, E. Gravitational waves from fallback accretion onto neutron stars. *The Astrophysical Journal*, 761(1):63, dec 2012. ISSN 0004-637X. doi: 10.1088/0004-637X/761/1/63.
- [86] Predoi, V. *Gravitational Waves and short gamma -Ray Bursts*. PhD thesis, 1993.
- [87] Privitera, S., Mohapatra, S. R. P., Ajith, P., et al. Improving the sensitivity of a search for coalescing binary black holes with nonprecessing spins in gravitational wave data. *Physical Review D*, 89(2):024003, jan 2014. ISSN 1550-7998. doi: 10.1103/PhysRevD.89.024003.
- [88] Ramadevi, Y., Sridevi, T., Poornima, B., and Kalyani, B. Segmentation And Object Recognition Using Edge Detection Techniques. *International Journal of*

BIBLIOGRAPHY

- Computer Science and Information Technology*, 2(6):153–161, dec 2010. ISSN 09754660. doi: 10.5121/ijcsit.2010.2614.
- [89] Roth, P. M. and Winter, M. Survey of Appearance-Based Methods for Object Recognition. *Transform*, (ICG-TR-01/08):66, 2008.
- [90] Saleh Al-amri, S. and Kalyankar, N. Image segmentation by using edge detection. *International Journal on Computer Science and Engineering*, 02 (03):804–807, 2010.
- [91] Santamaría, L. and Ott, C. D. Gravitational Wave Emission from Accretion Disk Instabilities – Analytic Models GWs from Accretion Disk Instabilities. 2013.
- [92] Saulson, P. R. *Fundamentals of Interferometric Gravitational Wave Detectors*. WORLD SCIENTIFIC, oct 1994. ISBN 978-981-02-1820-1. doi: 10.1142/2410.
- [93] Shapiro, B. and Lantz, B. LIGO-T1600467-v1: O1 Operating Mode Data. Technical report, 2016.
- [94] Singhal, P., Verma, A., and Garg, A. A study in finding effectiveness of Gaussian blur filter over bilateral filter in natural scenes for graph based image segmentation. In *2017 4th International Conference on Advanced Computing and Communication Systems (ICACCS)*, pages 1–6. IEEE, jan 2017. ISBN 978-1-5090-4559-4. doi: 10.1109/ICACCS.2017.8014612.
- [95] Sutton, P. J., Blackburn, L., Samarai, I. A., et al. X-Pipeline: an analysis package for autonomous gravitational-wave burst searches A comparison of methods for gravitational wave burst searches from LIGO and Virgo A first

BIBLIOGRAPHY

- [101] van Putten, M. H. P. M. Gravitational Waveforms of Kerr Black Holes Interacting with High-Density Matter. *The Astrophysical Journal*, 684(2):L91–L94, sep 2008. ISSN 0004-637X. doi: 10.1086/592216.
- [102] Van Putten, M. H. P. M. Gamma-ray bursts and gravitational radiation from black hole-torus systems.
- [103] Van Putten, M. H. P. M. Gravitational waveforms of Kerr black holes interacting with high-density matter. *The Astrophysical Journal*, 684:91–94, 2008.
- [104] Vikram Mutneja, D. and Mutneja, V. Methods of Image Edge Detection: A Review. *Journal of Electrical & Electronic Systems*, 04(02):1–5, jul 2015. ISSN 23320796. doi: 10.4172/2332-0796.1000150.
- [105] Waş, M., Bizouard, M.-A., Brisson, V., et al. On the background estimation by time slides in a network of gravitational wave detectors. *Classical and Quantum Gravity*, 27(1):015005, jan 2010. ISSN 0264-9381. doi: 10.1088/0264-9381/27/1/015005.
- [106] Winnemöller, H., Kyprianidis, J. E., and Olsen, S. C. XDoG: An eXtended difference-of-Gaussians compendium including advanced image stylization. In *Computers and Graphics (Pergamon)*, volume 36, pages 740–753. Pergamon, oct 2012. doi: 10.1016/j.cag.2012.03.004.
- [107] Zevin, M., Coughlin, S., Bahaadini, S., et al. Gravity Spy: integrating advanced LIGO detector characterization, machine learning, and citizen science. *Classical and Quantum Gravity*, 34(6):064003, mar 2017. ISSN 0264-9381. doi: 10.1088/1361-6382/aa5cea.