

Refining the Search for Sub-threshold Lensed Gravitational Waves

LIGO SURF 2021 Final Report

Storm Colloms,^{1, a}

Mentors: Alvin Li,^{2, b} and Alan J. Weinstein^{2, c}

¹*School of Physics and Astronomy, University of Edinburgh*

²*LIGO Laboratory, California Institute of Technology*

(Dated: October 24, 2021)

Gravitational lensing of gravitational waves is predicted in terms of strong, weak, and micro lensing, but has not yet been observed in any of these lensing regimes. Efforts to detect strongly lensed gravitational waves using data from Advanced LIGO and Advanced Virgo include searching for sub-threshold lensed images, which have been de-magnified and are undetectable with usual signal analysis methods due to a high noise background. In this work, I describe efforts to make improvements to targeted sub-threshold lensed gravitational wave search pipeline, which included investigating the waveform family used in the pipeline, as well as imposing conditions on the sky location of the lensed images. Spin-aligned and precession included waveform families were compared, and it was found that while the spin parameter of the gravitational wave event is so unconstrained, it cannot be said that either waveform family would incur loss of candidate images. Furthermore, a constraint was implemented on the sky localisation of lensed gravitational wave images, which can be approximated to come from the same sky location. By imposing this condition, the rankings of lensed image candidates with a stronger overlap to a target event were boosted as expected. This now needs to be implemented into the search pipeline for sub-threshold lensed images.

1. INTRODUCTION

In 2015 almost 100 years since they were first predicted [1], gravitational waves resulting from a pair of merging black holes were detected with the Advanced Laser Interferometer Gravitational-Wave Observatory (aLIGO) for the first time [2]. Since this detection, there have been on the order of 50 detection events over two and a half observing runs, with the number of detections increasing between each observing run as the instrument sensitivity and data analysis techniques are improved [3, 4].

Gravitational lensing has been studied extensively in the case of electromagnetic light [5, 6], as seen with the measured lensing of light from stars around the eclipsed sun [7], multiply imaged quasars and galaxies, and has been used to measure the distribution of matter in the universe as well as to constrain Hubble’s constant.

General relativity’s equivalence principle says this lensing should also occur for gravitational waves, yet this has not been observed despite many efforts. Gravitational lensing is the one of the fundamental properties of gravitational waves predicted by general relativity not yet to have been observed, others that have been observed include their speed, polarisation, and weak interaction with matter, as well as higher order modes.

Multiple efforts have been made so far to identify lensed gravitational waves within the current LIGO catalogues [8–15]. With the expected rate for lensing and the increasing gravitational wave detections, it is very likely that recent or near future data from LIGO will contain strongly lensed gravitational waves [16, 17]. This would allow not only

insight into comparisons of observations of lensing with general relativity, but also investigations into the structure of the universe and its matter distribution through the discovery of lens masses, such as probing the population of intermediate mass black holes [18].

Strong lensing creates multiple images of the lensed wave, and provided that the wavelength of the gravitational wave is much smaller than Schwarzschild radius of the lens, the magnification due to gravitational lensing will not depend on the gravitational wave’s frequency [19]. This means the wave profile of the lensed images would not be different from the wave if it were not lensed apart from an overall scaling factor, a time delay, and a phase shift. Some of these lensed signals may be de-magnified and thus considered “sub-threshold” waves, which would be buried in the noise background and therefore would have not previously detected.

Current searches for sub-threshold signals are ongoing, however there are still opportunities to improve these search pipelines to target them to sub-threshold lensed signals. This report details the methods used to refine the TESLA sub-threshold search pipeline [11], first investigating the waveform family used in the pipeline, and secondly implementing a condition of the sky location of the lensed gravitational wave. This constraint exploits the large uncertainty on current gravitational wave sky localisation compared to the much smaller lensed image separation, allowing the approximation that lensed images will come from approximately the same sky location. This would allow the ranking of triggers that have closer sky locations to a target event to be boosted, as these triggers would have been more likely to be lensed. If implemented into the search pipeline, this constraint would increase the detectability of sub-threshold signals and hope-

fully speed up the discovery of lensed gravitational waves.

This report is structured as follows: the background section details the properties of gravitational waves, the methods of detection and finding weak signals in noisy data, lensing of electromagnetic light and gravitational waves, and what efforts have been made so far to detect lensing from gravitational wave data. Section 3 then details the aims of the project. Section 4 describes the first portion of work on the project, investigating the waveform family used in the targeted sub-threshold search pipeline, detailing method, results and conclusions. Section 5 then details the method, results and conclusions of the second stage of the project, refining the search pipeline by considering the sky localisation overlap of lensed images. Section 6 then gives an overall conclusion to the project and insights on the future direction of this work.

2. BACKGROUND

Gravitational Waves

Gravitational waves were theorised to be self-propagating ripples in spacetime caused by accelerating masses, predicted by Einstein's theory of general relativity in 1916 [1]. The weak-field gravitational equation (equation 1) gives solutions of the form in equation 2 in the absence of mass, the form of these evidently hints at gravitational waves which stretch and squeeze the fabric of spacetime as they travel:

$$\square h_{ij} = -\frac{16\pi G}{c^4} T_{ij}, \quad (1)$$

$$h_{ij}(\vec{x}, t) \propto \exp(i\vec{k} \cdot \vec{x} - kct) \quad (2)$$

where G is the gravitational constant, c is the speed of light in a vacuum, T_{ij} is the stress-energy tensor, \vec{k} is the wave vector, \vec{x} the 3D space vector, t time, and i and j the spacial cartesian coordinates of the strain tensor h_{ij} . These solutions predict the speed of these waves as c , which has been verified with observed gravitational waves to be true to within -3×10^{-15} and $+7 \times 10^{-16}$ times the speed of light [20]. The form of the spacetime strain can also be expressed in terms of its Fourier transform from the frequency domain \tilde{h}_A , as observed at a detector:

$$h_{ij}(t, \vec{r}) = \sum_{A=+, \times} e_{ij}^A(\hat{n}) \int_{-\infty}^{+\infty} \tilde{h}_A(f) e^{-i2\pi f(t - \frac{\hat{n} \cdot \vec{r}}{c})} \quad (3)$$

where the gravitational wave strain is split across two polarisations, h-plus (h_+) and h-cross (h_\times). As a gravitational wave passes through a ring of test masses, the ring can be seen to be deformed as in figure 1. By gauge freedom, under the transverse-traceless gauge the spacetime strain tensor h only has spatial components without time, so the indices only run over space, which confines the deformations to the x-y plane where the wave propagates in the z direction. Therefore the

tensors $e_{ij}(\hat{n})$ are the polarisation tensors for the wave. In the long-wavelength approximation that is applicable for the LIGO detectors, the retarded time $t - \frac{\hat{n} \cdot \vec{r}}{c}$ can be simplified to simply t as the gravitational wave wavelength is much longer than the detector arm length \vec{r} . The resultant spacetime strain h_{ij} is a linear combination of these polarisations, giving the form of the deformations to be linear, circular, or elliptical.




Figure 1. A ring of free-falling test masses as deformed by h_+ (above) and h_\times (below) polarisations of a gravitational wave throughout one period of oscillations. It can be seen that the deformation from the h_\times polarisation is at a 45 degree angle to the h_+ polarisation. Figure From [21].

Detection of Gravitational Waves

Gravitational waves were first detected by aLIGO (Advanced Laser Interferometer Gravitational-Wave Observatory) on the 14th of September 2015 [22]. LIGO uses a Michelson-Morley interferometer, which measures the strain of the waves due to differences in the perpendicular arm lengths of the detector. A laser beam is split into the two arms by a beam splitter, the laser light then travels down the arms of length 4km and is reflected off mirrors at the end of each arm. To increase the path length to that which is needed for the required frequencies the light is reflected multiple times up and down the arms with Fabry-Perot cavities. Light recombines at a detector to usually produce destructive interference in the absence of noise or gravitational wave strain. In the presence of a gravitational wave the arms of the detector will change in length, the difference in arm length given by the strain h_{ij} times the arm length. Under the transverse-traceless gauge, the mirrors can be treated as free-falling, meaning the length of the arms does not change under a the wave strain and instead the time taken for the light to travel down the arm oscillates, thus the phase difference between the light at recombination will also oscillate. This changing phase difference is then used to measure the strain of the signal.

There are two LIGO detectors in Hanford and Livingston in the USA, and a third detector, Virgo in Italy. The multiple detectors can be used to confirm signals across the three locations, providing the time difference between the events is

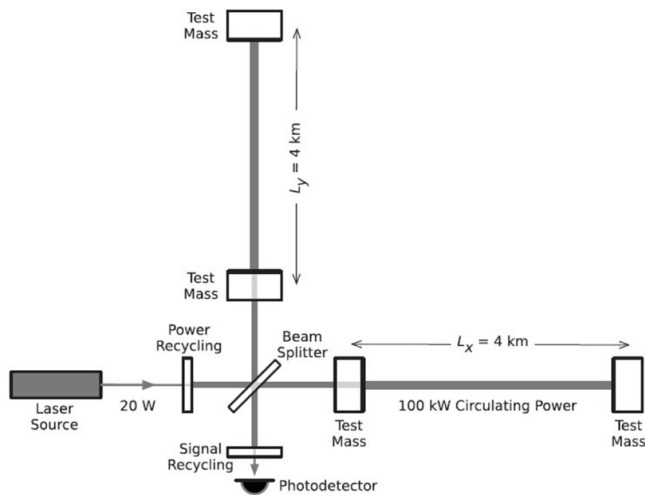


Figure 2. A simplified diagram of the Michelson Morley interferometer located at the LIGO and Virgo sites. 4km arms light reflected multiple times in Fabry-Perot cavities. The light recombines at the detector, which is set to destructive interference under the absence of a gravitational wave signal or terrestrial noise. Figure from [23].

at most the distance separating them divided by the speed of the waves, as well as strain the sky location of the event. If the same signal is being observed in more than one detector within a time window of the speed of light times the distances between the two detectors, the likelihood of the event can be boosted. This time window is at maximum $\pm 10ms$ between Hanford and Livingston detectors, and $\pm 25ms$ between Hanford and Virgo or Livingston and Virgo. Such criterion is known as the time coincidence criterion, and can strengthen the confidence that the event is a true signal.

Ground based detectors are tuned to detect gravitational wave frequencies in range from ~ 20 hertz to several kilohertz. Abundance of noise in the detectors makes it challenging to detect signals outwith this range, where there is a dip in the median noise frequency distribution, as seen in figure 3. The current sensitivity of the detectors allows gravitational noise strain to be detected down to strains of order $10^{-23} 1/\sqrt{Hz}$ around frequencies of 100 Hz.

The sensitive frequency band of these detectors corresponds to the typical frequency ranges of the inspiral, merger, and ringdown phases of compact binary coalescences (CBCs) in the stellar mass range, which can be categorised into merging events of binary black holes (BBH), binary neutron stars (BNS), or one of each (NSBH). The gravitational waves in the inspiral phase carry away energy, reducing the separation of the objects and their angular momentum. The merger phase occurs at the peak amplitude of the event in the time domain as seen in figure 4, when the sources coalesce into one body. Finally the body will release further gravitational waves while stabilising from the merging event. These three




Figure 3. The spectral frequency distributions of strain noise in LIGO and Virgo as of 2019. The lower frequencies are increasingly dominated by seismic and ground noise, while quantum noise dominates at higher frequencies. A so-called bucket can be seen around 100Hz, where current observations occur. Figure from [4]

stages, inspiral, merger, and ringdown, together constitute a coalescence and are shown in figure 4.

The first event at aLIGO on the 14th of September 2015 was the merging event of two stellar mass black holes (GW150914) [23], and since then approximately 50 further mergers have been identified from neutron stars around several solar masses large up to black holes as massive as several hundreds of solar-masses [3][4].

Other predicted sources of gravitational waves include a stochastic background of unresolved gravitational wave signals at a further distance than current CBC detections; continuous waves which could be propagating from rotating neutron stars that are not perfectly spherical; and burst gravitational waves which encapsulates waves from poorly modelled sources [24].

While these mergers are some of spacetime's most violent events, the gravitational waves spread across spacetime until they are detected on Earth. The amplitude of the strain of these gravitational waves is thus reduced from the time of the event, and then measured as the fractional change in arm length in the detectors, i.e. the ratio of spacetime distortion. For the current observations, this strain is of order $h(t) = D^{ij}h_{ij}(t) \sim 10^{-21}$ or smaller, where D^{ij} is a tensor that depends on the geometry of the detector.

To be able to detect the strain of gravitational waves above




Figure 4. The strain of gravitational wave GW150914 versus time alongside a graphic depiction of the inspiral, merger, and ringdown phases (top). The bottom panel shows the black hole relative velocity, which increases exponentially until the merger phase, and the black hole separation which decreases until the same point. Figure from [23].

noise in the detector, efforts have been made to improve the detector sensitivity and reduce the noise. Noise sources include quantum, seismic, and thermal noise over large frequency bands, as well as fixed frequencies of noise such as 60Hz mains power lines which constitute the spectral lines in figure 3. Noise such as seismic noise from earthquakes can be identified and removed by comparing real time data from seismometers and microphones to the strain data received at the laser detector. Despite efforts to remove many of the noise sources, strain data from gravitational waves still remains “hidden” behind the noisy data, as seen figure 5. The main aim of gravitational wave signal analysis is then to be able to distinguish the gravitational wave from the background noise.

Finding Weak Signals in Noisy Data

Despite efforts to reduce noise, the noise is still much larger in magnitude than the typical CBC gravitational wave in the LIGO frequency band. Therefore, a correlation between the strain data in the detector and the approximate known shape of the target waveform is found. This is called matched filtering. The shape of the waveform can be determined by 15 parameters of the coalescence; 8 intrinsic parameters: the mass of each coalescing body m_1 and m_2 , their spins \vec{S}_1 and \vec{S}_2 ; and 7




Figure 5. The strain data versus time for gravitational wave GW151226 in the Hanford and Livingston detectors, filtered with a 30–600 Hz bandpass filter to remove large fluctuations in strain outside this frequency range, shown in blue and red. Shown in black is the waveform of the gravitational wave of the strain data for the wave with best fitting parameters to the data. It can be seen that the noise is much greater than the wave strain, thus matched filtering is needed to extract the waveform. Figure from [25].

extrinsic parameters: the time of coalescence t_c , the reference phase φ_c , the sky position in right ascension and declination α and δ , luminosity distance d_L , and the orbital orientation Θ_{Jn} , ψ . Additionally neutron stars can be characterised by another parameter, tidal deformability $\tilde{\Lambda}$ for the total of 16 parameters.

In order to find a gravitational wave event, the signal-to-noise ratio (SNR), a measure of the strength of the signal assuming it matches a given template, can be maximised for a range of templates, hence finding the best fitting template. While the parameters that determine the waveforms of the templates are continuous, the templates used to perform matched filtering are discrete and the collection of templates is named the template bank.

The data received at the detector $s(t)$ is given by the addition of the noise signal $n(t)$ and the gravitational wave signal $h(t)$:

$$s(t) = n(t) + h(t). \quad (4)$$

For a given filter with filter function $P(t)$, which gives the whitened function of the template, the matched filtering output can be given by:

$$\hat{s} = \int s(t)P(t)dt. \quad (5)$$

The SNR can then be defined as the expectation value of the matched filtering output when a signal is present, $\langle S \rangle$ divided by the root mean square of the output when there is no GW signal, N :

$$\rho = \frac{\langle S \rangle}{N}, \quad (6)$$

where $\langle S \rangle$ can be found with:

$$\langle S \rangle = \int \langle s(t)P(t) \rangle dt \quad (7)$$

$$= \int \langle (n(t) + h(t))P(t) \rangle dt \quad (8)$$

$$= \int \langle h(t)P(t) \rangle dt \quad (9)$$

$$= \int \tilde{h}(f)\tilde{P}^*(f)df. \quad (10)$$

Here $\langle n(t) \rangle$ and thus $\langle n(t)P(t) \rangle$ can be taken as zero assuming the noise is both Gaussian and stationary. This is can be assumed to be the case under the assumption that no glitches, non-astronomical disruptions in the Gaussian and stationary nature of the noise, occur.

N is then:

$$N^2 = \langle \hat{s}^2 \rangle - \langle \hat{s} \rangle^2 \quad (11)$$

$$= \langle \hat{s}^2 \rangle \quad (12)$$

$$= \iint P(t)P(t')\langle n(t)n(t') \rangle dt dt' \quad (13)$$

$$= \frac{1}{2} \int S_n(f)|\tilde{P}^*(f)|^2 df, \quad (14)$$

where S_n is the power spectral density of the noise, or the power distribution in the noise among all frequency components.

If a signal then matches with a template to create a SNR surpassing a certain threshold at a certain time in the data, it is then called a trigger. SNR based ranking of signals is a good starting point, however it cannot be used solely as it assumes that the noise is Gaussian and stationary. In the case of noise glitching and non-stationarity, a false-alarm rate (FAR), which signifies the rate that real detector noise, including glitches, could produce the trigger under consideration, amongst other ranking statistics is used [26, 27].

The GstLAL Search Pipeline

GstLAL is one of the pipelines for matched filtering and data reduction of the gravitational wave data [27]. The pipeline can be performed either offline once all of the data has been recorded, or low-latency, i.e. with minimal delay from the recording of the data. Low-latency methods allow for rapid electromagnetic follow-up of any candidate gravitational waves. The primary steps in the offline mode of the GstLAL pipeline are shown in figure 6.

After the data have been received and recorded from the 3 detectors, the first step in the pipeline is to estimate the PSD for the chunk of data. Above, the PSD was found directly

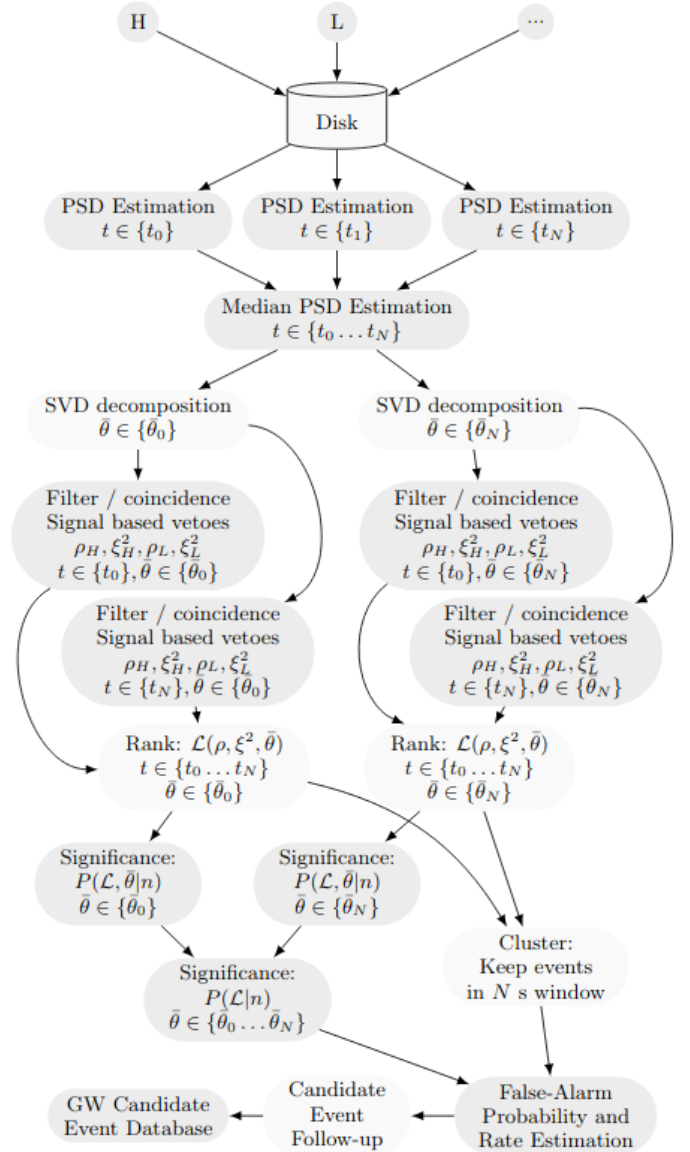


Figure 6. Diagram of the offline search mode of the GstLAL pipeline, following data retrieval, PSD estimation, filtering and ranking, and follow-up parameter estimation. Figure from [27]

from the equation:

$$\langle \tilde{n}(f)\tilde{n}^*(f') \rangle = \frac{1}{2}S_n(f)\delta(f-f'). \quad (15)$$

However this cannot be used for real data as it may contain an actual signal and not just noise, and will also contain glitches which would make the noise non-Gaussian non-stationary. Therefore the data is split into small time slices and Fourier transformed into the frequency domain, from which the median of these is taken to to be the median PSD for said chunk of data, reducing the influence of a possible signal on the noise estimation [27].

Obvious glitches are then removed from the data, by firstly whitening data series by dividing it by the square root of the PSD. The amplitude of the data is scaled so that variance of the data amplitude is then 1. If there is any point in the whitened data that has a variance exceeding a certain value, typically 50 or 100, it can be said to be a glitch and can be removed by zeroing the data around that point to 0.25s either side, in a process known as gating. This removes obvious glitches that would create high SNR triggers in the data. Alternatively, if there is a known source of glitched data from instrumental and environmental monitors at the detectors, such as an earthquake, this can be flagged so the data can be vetoed (set to zero) around this time.

Since the template bank is discrete yet the parameters of gravitational wave signals are continuous, if a signal occurs that does not exactly match a filter this can incur a loss of SNR. In matched filtering, this can be mitigated by increasing the number of templates. To ensure there is not a loss in latency, templates are reduced to vectors which match part of the signal. This is known done as the LLOID process, which encompasses chopping templates into different time slices, down sampling each slice, and single-value decomposition (SVD). The total SNR for each trigger is then found from the total of the SNR from the top select eigenvectors.

In addition to SNR, each trigger must be categorised by the difference between the data and the best-fit matched filter template, ξ^2 . This allows for further distinction between glitches and true signals, as a high SNR trigger may not truly match well to the filter in other coalescence phases outwith the high-amplitude and thus high SNR merger phase.

The SNR (ρ) and the ξ^2 of the triggers, as well as the filter parameters $\bar{\theta}$ can be used to determine the likelihood ratio \mathcal{L} , given by equation 16:

$$\mathcal{L}(\rho_1, \xi_1^2, \dots, \rho_d, \xi_d^2, \bar{\theta}) = \frac{P(\rho_1, \xi_1^2, \dots, \rho_d, \xi_d^2, \bar{\theta} | s)}{P(\rho_1, \xi_1^2, \dots, \rho_d, \xi_d^2, \bar{\theta} | n)}, \quad (16)$$

for number of detectors d , and where $P(\dots | s)$ is the probability of observing the trigger (...) given a signal, and $P(\dots | n)$ is the probability of observing (...) given noise.

The false alarm probability, $P(\mathcal{L} | n)$, can also be calculated as the probability of measuring a given \mathcal{L} if the data contains only noise. From this, the probability of obtaining a trigger of log-likelihood ratio greater than a certain threshold \mathcal{L}^* may be found by integrating the false alarm probability from this threshold to infinity. The false alarm rate (FAR), defined as how often the noise is expected to yield an event with a log likelihood greater than this threshold, can be calculated by dividing this cumulative probability by the observing time:

$$FAR = \frac{\int_{\mathcal{L}^*}^{\infty} d\mathcal{L} P(\mathcal{L} | n)}{T}, \quad (17)$$

where T is the observing time.

In the GstLAL search pipeline, if the trigger is detected in more than one detector then its ranking is promoted, however single detector triggers can still be ranked highly while incurring a penalty.

Following the final ranked list of triggers, the highest ranking triggers undergo follow-up parameter estimation using Bayesian analysis. The parameters estimated from matched filtering are just those best fitting to the template, and the template parameters could actually differ considerably from the true event parameters especially in areas of the parameter space with lower template density, and thus parameter estimation gives more reliable results for 15 (16 for a neutron star) parameters of the CBC.



Gravitational Lensing of Gravitational Waves

Another consequence of the bending of spacetime is that of gravitational lensing. As light travels on the shortest path through spacetime, it's path may be bent by curved spacetime. This can cause the apparent position of the bent light source to change according to an observer on the other side of the lens, as seen in figure 7.




Figure 7. The bending of radiation around a massive object, here the sun, to reach an observer on Earth who would presume the light has travelled in a straight line, hence altering the source objects apparent position from its true position. Figure from [5].

Gravitational lensing was first observed in electromagnetic light, which came to the scientific forefront in the early 20th century, alongside general relativity. The first observations of gravitational lensing were seen in the measured lensing of light from stars around the eclipsed sun. Since then, the presence of multiple strong lensed images of electromagnetic light, as well as Einstein rings, were found. Efforts have been made to identify these lensed images, such as lensed Quasars, measuring the images, redshifts, spectra, variability, and presence of a lens mass to confirm that they are truly lensed [5, 6].

In addition to strong lensing of a source, which creates multiple images, a source may also be weakly lensed, where the path of the emitted electromagnetic light is altered slightly by intervening matter, and the shape of the source object changed slightly. Furthermore, microlensing can be observed as a time-dependant lensing event.

Strong lensing also effects the arrival time of the light due to the difference in path lengths for different paths, where the speed of light is the same in all frames, including non-inertial frames like curved space. This creates different images for the observer, and has been observed in variable quasars, where multiply imaged quasars display the same variability with each image having a distinct time delay [28].

Lensing has been used in a variety of useful ways, including constraining Hubble's constant by measuring the time delay between variabilities in images. Statistically analysing lensed objects can give insight in the population and distribution of these lensed sources, and hence the cosmological model of the universe [6].

Under the equivalence principle, gravitational lensing should also occur for gravitational waves. While lensing of light may be obscured due to dust between the lensed source and the observer, this is not a problem for gravitational waves which are not impacted by the obscuring matter. This means that the observations of lensed gravitational waves could allow further probing of the structure of the universe that cannot be gathered from the lensing of light.

The strain of the lensed gravitational wave can be related to the strain of the unlensed gravitational wave:

$$h_{+, \times}^{lensed}(f) = F(f)h_{+, \times}^{original}, \quad (18)$$

where $F(f)$ is the amplification function that can depend on the frequency of the wave, as well as parameters such as the lens mass and the distances between the observer, the lens and the source [19]. Two simple models can give two forms of this equation, taking the lens as a point-mass, or a singular isothermal ellipsoid (SIE), although these are vastly simplified and not comparable to real astrophysical lenses. These models create two or four images of the lensed wave, with different amplifications and time delays. However, many models exist to approximate gravitational lensing, many being more complicated and resulting in different numbers of images amongst other different lensing effects.

Assuming the geometric optics limit holds such that the wavelength of the gravitational waves is much smaller than the Schwarzschild radius of the lens, $F(f)$ can be shown to be independent of the frequency of the wave, instead solely modifying the wave's amplitude and phase [19]. This is also known as the lensing being achromatic. This limit holds in the cases with stellar mass coalescing source objects being lensed

by massive objects such as supermassive black holes or star clusters, but breaks down in the case of merging events from supermassive black holes lensed by similarly massive objects, or when stellar mass sources are lensed by intermediate black holes of mass $\sim 10^3 M_{\odot}$. However, for the LIGO frequency band, it is expected that any lensed events observed would be able to approximated by the geometric optics limit. The lensing equation dictating this case would then be of the form:

$$h_j^{lensed}(f, \bar{\Theta}, \mu_j, \Delta t_j, \Delta \phi_j) = F(f) \times h^{original}(f, \bar{\Theta}, \Delta t_j), \quad (19)$$

where

$$F(f) = \sqrt{|\mu_j|} \times \exp(i \text{sign}(f) \Delta \phi_j), \quad (20)$$

and $\bar{\Theta}$ are the parameters of the gravitational wave, $\sqrt{|\mu_j|}$ is the magnification factor, Δt_j is the time delay between a pair of lensed images, and the exponential term is the morse phase shift, which depends on whether the image is of type 1, 2, or 3 [14]. For a type 1 image, the phase shift $\Delta \phi_j$ would be equal to 0, for type 2 equal to $-\pi/2$, and $-\pi$ for type 3. Received at the detector, this amplitude modification scales the SNR of the lensed wave, as seen in figure 8. For some SNR threshold that needs to be exceeded for an event to be classified as a trigger, some lensed images may be de-magnified and not exceed the trigger threshold. These are called sub-threshold events [10, 11].

For the case of two lensed images, The probability distribution of the amplification factor has been found to differ between each image, one being very likely to be demagnified while the other magnified to typically a factor of 4 [12]. The typical time delay between these images would also be on the order of minutes, although could be as long as weeks, or months, depending on the lens mass [12, 16, 17].

The latest analysis of O3a search run data sets an expected lensing rate on the order of $10^{-3} - 10^{-4}$, as in one in every hundred to thousand waves could be lensed, depending on the lens type and whether a single or double image is detected [8]. Another study predicts that 30% of lensed waves in the strong lensing regime would give 4 gravitational wave images [17]. These statistics make it highly likely that LIGO data will soon contain or has already contained evidence of lensed waves, giving strong motivation to pursue this search.

Searches So Far

Search runs O1 (2015-2016), O2 (2016-2017), and O3a (2019) have been used previously to look for evidence of lensed gravitational waves through a variety of methods.

In the first two observing runs, data were analysed to look for three forms of lensing evidence: gravitational wave magnifications altering the expected redshift and chirp mass distribution of the detected BBHs; super-threshold pairs




Figure 8. Waveforms of lensed counterparts plotted with waveform template family IMRPhenomPv2. The gravitational wave source is a CBC of two masses of $10 M_{\odot}$ and $20 M_{\odot}$, spins of 0.9 and 0, at luminosity distance $1 Mpc$. The inspiral stage of these counterparts would actually last much longer than 0.2s, but the start of the waveform is tapered in at this time to avoid overlapping waveforms of multiple events, extra computational work in producing such long data, and having an abrupt start which would produce Gibbs noise in any Fourier transforms. The dimensionless source position relative to the central axis passing through the observer and the lens is 0.5, giving the ratio of magnifications to be ~ 3 assuming a point mass lens; and the time delay for a lens of mass $2.5 \times 10^4 M_{\odot}$ is as shown 0.5 seconds [19]. The magnified counterpart arrives before the de-magnified wave, which is shorter than the expected delay for lensed images of gravitational waves. It can be seen that the overall amplitude of the lensed waves are modified by the lensing but their waveform is identical. If these waves were captured in noise and the SNR versus time were plotted, the first magnified wave might have high enough SNR to surpass the trigger threshold, with the de-magnified counterpart could be buried, making it sub-threshold.

of gravitational waves amongst already confirmed events; and instances of wave optics effects in the waveform of the signals where the geometric optics limit did not hold [9]. No compelling evidence for lensed gravitational waves was identified with these methods.

Furthermore, multiple techniques were used to search for sub-threshold lensed images starting in O1 and O2 search runs. The results from [10] show no evidence of lensed waves, while those from [11] still require parameter estimation to claim the candidates as lensed waves.

These methods were further developed for O3a lensing analysis, which looked at previous lensing rate with O3a detector sensitivities; and for multiple images in posterior-overlap and joint parameter estimation analysis, as well as sub-threshold lensing analysis and searches for microlensing [8]. Again there was no compelling evidence for lensed gravitational waves.

The sub-threshold waves are considerably smaller than the noise-background and thus require further analysis to extract them. One such analysis pipeline is named TESLA (TargetEd Subthreshold Lensing seArch), its workflow detailed in 9.

TESLA searches for sub-threshold signals by first creating a set of injections that are scaled versions of a super-threshold event. The SNR of an event is inversely proportional to luminosity distance, thus new injections can be created by scaling existing templates to be at larger distances, reducing their amplitude and hence SNR without otherwise redshifting the waveform, mirroring the effect of the lensing de-magnification. Care is taken as to specify the source frame masses when generating the scaled waveforms, as changing the luminosity distance while preserving the detector frame masses would change the source frame masses, which need to be the same as those of the original super-threshold event. These scaled signals are injected into the data, and then a GstLAL search run is performed, and the matched filter templates able to found the injections are identified, as these templates may also be able to identify sub-threshold signals of the corresponding target super-threshold event.

These templates are then used as a reduced template bank, which can be used to perform a new targeted search run. Decreasing the templates in the template bank reduces the noise background, without reducing the ability to detect the gravitational wave signal. Any candidates from this targeted search are then followed up with Bayesian analysis to perform parameter estimation [29].

3. OBJECTIVES

This project aims to refine the TESLA pipeline in order to improve its efficiency and in turn the chances of finding evidence of lensing of gravitational waves. Two areas of the pipeline were investigated: the waveform family used in the pipeline to evaluate the templates used in matched filtering, methods and results are discussed in section 4; and imposing a condition on the sky location of the lensed images, methods and results of this are discussed in section 5.

The waveform template family used in the TESLA pipeline for creating the matched filtering templates in identifying sub-threshold gravitational waves is the spin-aligned approximant IMR-PhenomD. However, using a different waveform family may alter the reduced template bank and any possible candidates. The difference between spin-aligned and precession included families, and the effect of the choice of either family on the pipeline was investigated in order to find if either waveform family made a significant difference on the resulting lensed candidates.

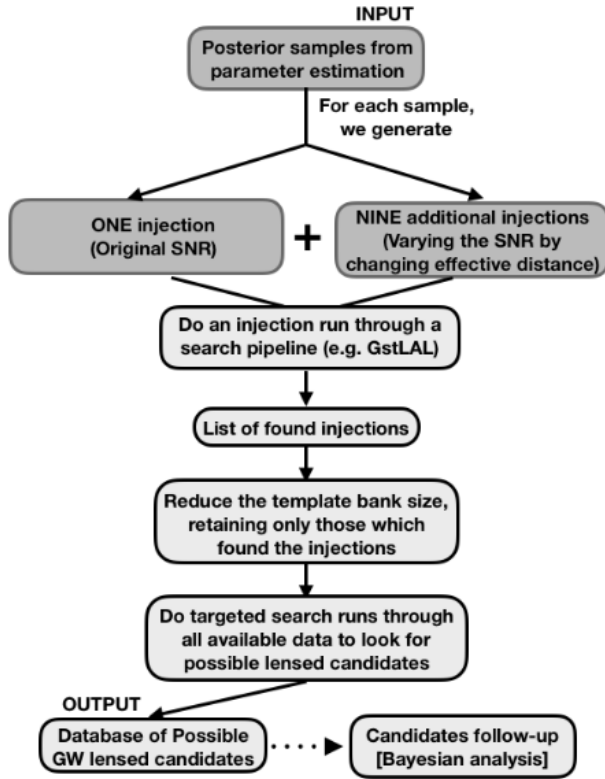


Figure 9. The TESLA pipeline to search for sub-threshold gravitational waves as used in [11]. The template bank is reduced by injecting signals with altered and original SNR into the data and only preserving the templates which recover these injections. The new template bank is used to search through all the data, which now has a reduced noise background due to the reduced template bank. Follow-up then needs to be performed on the candidates from this targeted search. Figure from [11].

Furthermore, given that in current analyses the sky localisation of gravitational wave signals is constrained to up to thousands of degrees squared, while lensed image separation would be on the order of arcseconds, it can be reliably stated that lensed gravitational waves will come from approximately the same sky location [4]. This property can be exploited to place another constraint on the search pipeline, aiming to give lower ranking statistics to candidate images of wide separation, and boost the rankings of candidates with higher sky localization overlap.

4. CONSIDERING THE EFFECT OF WAVEFORM FAMILY ON THE TARGETED SUB-THRESHOLD SEARCH PIPELINE

It can be argued that the a bad fitting waveform family used to calculate the strain evolution of a gravitational wave could lead to lensed signals being missed in the pipeline. One of the main differences between these waveform families is whether they consider precession. Spin-aligned waveforms take the

spin of the merging objects to be in the same direction as the orbital angular momentum, while precessing waveforms take any spin orientation to calculate the strain of the gravitational waves. Two parameters are used to define the in-plane spin and spin-aligned components of the spin, named χ_P and χ_{eff} respectively and defined as:

$$\chi_P = \max[\chi_1 \sin\Theta_1, (\frac{4q+3}{4+3q})q\chi_2 \sin\Theta_2], \quad (21)$$

$$\chi_{eff} = \frac{\chi_1 \cos\Theta_1 + q\chi_2 \cos\Theta_2}{1+q}, \quad (22)$$

where χ_1 and χ_2 are the dimensionless component spins of both objects, Θ_1 and Θ_2 are the tilt angles between the component spins and the orbital angular momentum, and q is the mass ratio of the merger.

All previously recorded super-threshold signals for O1, O2, and O3a have already recovered from previous lensing searches using both precessing and spin-aligned waveform families within the search pipelines [8]. However, it is still necessary to investigate whether the sub-threshold signals would be effectively recovered with the same SNR using either waveform.

To do this, I found the match between precessing and spin-aligned waveforms of previous gravitational wave events from O1 and O2. The BNS event GW170817 was disregarded as the lensing searches will primarily target BBH merger events, as they will be louder farther away and thus are more likely to have lensing masses along the line of sight from the event to Earth.

Firstly, the maximum posterior was found for each of the 15 parameters of the merger, where the posterior samples were found from [30], and the maximum likelihood posterior sample found to calculate the match, as these samples are the ones that best fit the data. Two sets of these posteriors were found for each event, one for spin-aligned posteriors and one for effective precession posteriors.

The antenna pattern function was also considered in order to calculate the actual strain response in the detector. The detector response for each polarisation and each detector d , $F_{+,d}$ and $F_{\times,d}$, were combined with the strain for each polarisation, h_+ and h_\times to give the total antenna response for each detector h_{tot} :

$$h_{tot,d} = F_{+,d} \times h_+ + F_{\times,d} \times h_\times. \quad (23)$$

The waveforms were then generated in each detector for each h_{tot} using both a spin-aligned waveform family, SEOBNRv2 and an effective precession family, IMRPhenomPv2, as shown for GW170104 in figure 10. The

matches between these two waveforms found with `pycbc` and the average results between detectors from these matches found and plotted in figure 11.




Figure 10. The time domain waveforms for GW170104 as seen in the Hanford detector, calculated using spin-aligned waveform family SEOBNRv2 in blue and precession-included waveform family IMRPhenomPv2 in orange. The precessing waveform displays obvious signs of amplitude modulation. The match between these two waveforms was 58%.

It was found that while many of the events were above 90% match between waveforms, GW170104 exhibited a lower average match of 54%. It is possible that this merger was therefore precessing. However, the posterior distribution for the in-plane spin components of each body χ_P was found to be very wide, and therefore it cannot be determined if the orbit is precessing. Therefore may not necessarily be SNR loss between the two waveform families. However, if the spin parameter χ_P were better constrained and the system was definitely precessing, then the precessing waveform would be truer match to the signal, and thus the choice of a spin-align waveform family could incur a 46% loss in SNR for the match of 54%. Thus it would be important to choose a precessing waveform in this case. However, with current methods the spin parameter is not often measured this precisely, and so the choice of waveform family cannot be said to currently alter the detectability of sub-threshold gravitational waves.

5. BOOSTING RANKINGS OF SIMILAR SKY LOCATION CANDIDATE PAIRS

In `GstLAL`, the likelihood ratio depends on the differences in coalescence phase $\Delta\phi$ and time Δt between detectors [26]. The expression in equation 16 can be altered to explicitly include the terms $\Delta\phi$ and Δt which describe the coincidence criteria, such that for each set of two detectors, it can be defined as:





Figure 11. The average match from both Hanford and Livingston detectors between spin-aligned and precessing waveform families. Eight of the events have a $> 90\%$ match, although GW170104 displays a match of approximately 53%. This event could be precessing and therefore not well defined by a spin-aligned waveform. This could cause problems in the sub-threshold lensing search if a spin-aligned waveform is used to search for a precessing sub-threshold signal, as it may incur heavy SNR loss.



$$\mathcal{L} = \frac{P(\vec{D}, \vec{d}, \vec{\rho}, \vec{\xi}^2, \Delta\phi, \Delta t | \text{signal})}{P(\vec{D}, \vec{d}, \vec{\rho}, \vec{\xi}^2, \Delta\phi, \Delta t | \text{noise})}, \quad (24)$$

where \vec{D} is the pair of horizon distances $\{D_1, D_2\}$ to each detector $\{d_1, d_2\}$ in \vec{d} , $\vec{\rho} = \{\rho_1, \rho_2\}$ are the SNRs and $\vec{\xi}^2 = \{\xi_1^2, \xi_2^2\}$ the consistency checks in each detector. The consistency check ξ_i^2 is important to distinguish between high SNR glitches that have a large residual between the glitch waveform and the template versus a well-matching signal.

Δt and $\Delta\phi$ depend on the sky location of the source and the positions of the detectors, as a wave coming from a different direction relative to the line between the pairs of detectors will have different times of arrival. Any difference in these times will introduce phase difference between the waves at each detector in the pair. The horizon distance is how far the detectors are each able to see, depending on their sensitivity at each point in time, and thus alter the log likelihoods, while

the set of detectors operable at the time of detection will alter the log likelihood as more detectors operable will be able to constrain sky localisation and contribute to the coincidence criterion.

For an injection run performed covering the full BBH and BNS parameter space, the probability density function (PDF) can be found for each $\Delta\phi$ and Δt [26]. This is shown in figure 12 for the Hanford and Livingston detectors, constrained to $|\Delta t| > 10\text{ms}$ as this is the maximum time of travel at the speed of light between these two detectors. Thus, the ranking statistic \mathcal{L} is modified to boost rankings of high probability $\Delta\phi - \Delta t$ values, which would be the green areas in figure 12.




Figure 12. The probability distribution function $P(\Delta\phi, \Delta t | \rho_{network})$ across $\Delta\phi$ and Δt parameter space, $\rho_{network} = \sqrt{\rho_1^2 + \rho_2^2}$ and $\rho_{H1} = \rho_{L1} = 10$.

In contrast to a non-lensing search run where the signal is assumed to come from any sky location, lensed image searches can be targeted to a particular sky location as lensed image separation is expected to be on the order of arcseconds. Comparing this to the wide uncertainty of gravitational wave sky localisation which can be on the order of thousands of degrees squared, it can be approximated that in the search pipeline lensed images will come from the sky location. Thus the sub-threshold search pipeline can be further targeted to assume a lensed image candidate will come from the same sky location as that of a super-threshold target event.

For a targeted search run, constraining the sky location of the PDF calculation to that of a super-threshold event alters the PDF of $\Delta\phi$, changing the shape of the PDF and decreasing the high probability area of the parameter space. Therefore this constraint can be built in to the current GstLAL TESLA pipeline through the introduction of the targeted PDF, to alter the log likelihood ratio and boost the ranking of signals which

are closer in sky location to the target, and thus refine the search for lensed gravitational waves.

For a given set of detectors, sensitivity, and signal, the probability as shown in the PDF figure above for the SNR, phase and time difference between detectors obeys the following equation [31]:

$$P(\vec{\rho}, \vec{\phi}, \vec{t} | \vec{d}, \vec{D}, signal) \propto P(\ln \vec{D}_{eff}, \vec{\Delta\phi}, \vec{\Delta t} | \vec{d}, signal) \times |\vec{\rho}_{network}|^{-4}, \quad (25)$$

where $\vec{\phi}$ and \vec{t} are the arrival phase and arrival time of the signal at each detector respectively, and $\rho_{network}$ is the SNRs in each detector added in quadrature. D_{eff} is the effective distance to the signal source and can be used to estimate the probability instead of the SNR due to the relationship:

$$\rho \propto \frac{1}{D_{eff}}. \quad (26)$$

The effective distance combines the source distance, detector antenna pattern and CBC inclination angle into one quantity, as in a single detector the location and orientation of the signal are not known, hence the single observable of effective distance is useful [32]. It is given by:

$$D_{eff} = D \left[F_+^2 \left(\frac{1 + \cos^2 \Theta_{Jn}}{2} \right)^2 + F_\times^2 \cos^2 \Theta_{Jn} \right]^{-1/2}, \quad (27)$$

where F_+ and F_\times are the antenna response function for each polarisation of the signal, and Θ_{Jn} is the inclination angle of the source.

The factor of $|\rho_{network}|^{-4}$ in equation 25 is due to the fact that the probability of observing a signal at a distance D over an area dD is proportional to $D^2 dD$, and the probability of observing a signal of some SNR ρ is then proportional to $\rho_{network}^{-4}$:

$$P(D) dD \propto D^2 dD \quad (28)$$

$$P(\rho) = P(D) \frac{dD}{d\rho} \propto \frac{1}{\rho^4}. \quad (29)$$

In GstLAL, the probability given on the right-hand side of equation 25 is evaluated with a further simplification which grids of points across the sky. For each point in this grid, the time, phase and effective distance at each tile point is calculated over every inclination angle. The difference between arrival time and arrival phase for each detector for each point on the sky is then calculated for some effective distance. The probability is then found considering these quantities, and for any D_{eff} the PDF of $\Delta\phi$ and Δt can be returned for each detector pair. The final probability then depends on this grid of

$\Delta\phi - \Delta t$ and a covariance matrix calculated within the pipeline [31]. First taking:

$$\vec{\lambda} \equiv \left\{ \Delta \ln \vec{D}_{\text{eff}}, \vec{\Delta t}, \vec{\Delta\phi} \right\}, \quad (30)$$

and:

$$\vec{\Delta\lambda}_i \equiv \vec{\lambda} - \vec{\lambda}_{\text{mi}}, \quad (31)$$

where $\vec{\lambda}_{\text{mi}}$ is a set of model vectors. Then for covariance matrix $\Sigma_{\vec{\lambda}}$ which can be decomposed into the matrices $\mathcal{C}\mathcal{C}^T$:

$$P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{\text{mi}}) = \frac{1}{\sqrt{(2\pi)^k |\Sigma_{\vec{\lambda}}|}} \exp \left[-\frac{1}{2} \vec{\Delta\lambda}_i \Sigma_{\vec{\lambda}}^{-1} \vec{\Delta\lambda}_i^T \right]. \quad (32)$$

Taking $\Delta\vec{x}_i = \vec{\Delta\lambda}_i \mathcal{C}$:

$$P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{\text{mi}}) \propto \exp \left[-\frac{1}{2} \Delta\vec{x}_i^2 \right]. \quad (33)$$

We have:

$$P(\vec{\lambda} | \vec{O}, s) \propto \sum_i P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{\text{mi}}) P(\vec{\lambda}_{\text{mi}}) \quad (34)$$

$$\propto \sum_i P(\vec{\lambda} | \vec{O}, s, \vec{\lambda}_{\text{mi}}) \quad (35)$$

$$= \sum_i \exp \left[-\frac{1}{2} \Delta\vec{x}_i^2 \right]. \quad (36)$$

To constrain the search by the sky location of a target event, the tiling across the sky was changed to only tile the area of the 90% credible region of said target event, as found by previous search results from GraceDB [33]. Each time phase and effective distance were only sampled over the pixel locations of this reduced credible area, and then the $\Delta\phi - \Delta t$ PDF was calculated as before for each detector pair in Hanford, Livingston, and Virgo. This resulted in PDFs with reduced high probability areas, as shown in figure 13 which gives the targeted S190408an event sky location in as seen in Hanford and Livingston, at the true events SNRs in each detector, as given by GraceDB [33]. The masses and spins from the search results were used in the pipeline to calculate the covariance matrix, from the same search results on GraceDB. It can be seen that the search template parameters for $\Delta\phi$ and Δt in Hanford and Livingston agree to a reasonable extent with the PDF as shown in figure 13. The offset of the parameters given by GraceDB and the peak of the PDF could be due to the used to calculate the covariance matrix masses not being accurately measured from the pipeline, as they are values from the best fitting matched template before posterior analysis.

The PDFs for each detector pair at each detector SNR equal to 4 were generated using this target event S190408an,

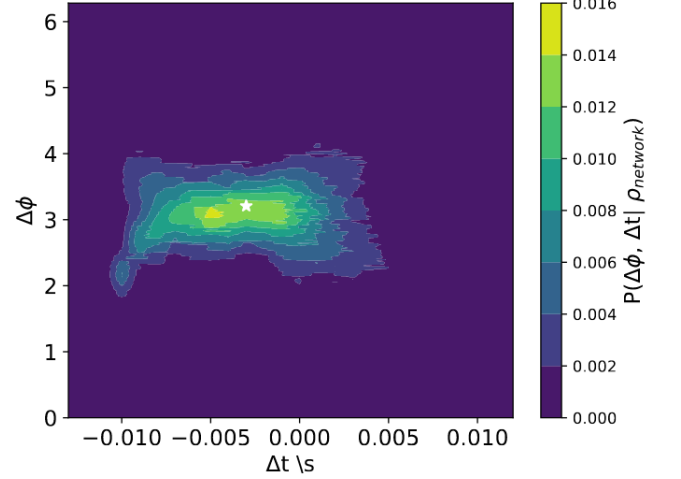


Figure 13. The probability distribution function $P(\Delta\phi, \Delta t | \rho_{\text{network}})$ across $\Delta\phi$ and Δt parameter space, targeted to the 90% credible area of super-threshold event S190408an, and at SNR in Hanford equal to 9.3975878, and 9.9423523 in Livingston, three true event parameters (from [33]). The white star shows the event's values of $\Delta\phi = 3.20841323$ radians and $\Delta t = -0.003$ s as also found from GraceDB.




Figure 14. The probability distribution function $P(\Delta\phi, \Delta t | \rho_{\text{network}})$ across $\Delta\phi$ and Δt parameter space in detectors Hanford and Livingston, targeted to the 90% credible area of super-threshold event S190408an, and at SNR of 4 in all detectors. This PDF was used to conduct a search run for O3a, and results shown in table I

as shown in figure 14. The low SNR was used in order to target low SNR sub-threshold signals. These targeted PDFs were then used in a GstLAL search run performed on entire O3a data.

From this targeted search run, the resulting log likelihood ratio, FAR, and overall ranking for the target event S190408an

as well as lensed candidate image pairs to the target event were compared to those of the previous search without a targeted sky location PDF. The results are shown in table I.

The target event’s ranking was decreased using the targeted PDF, which was not expected as the PDF was specifically targeted to its own sky location. However, this may be due to targeting the detector SNRs equal to 4, while the super-threshold event was detected at a higher SNRs in each of the Hanford and Livingston detectors of ≈ 9 , and thus did no longer correspond to a high probability are on the lower SNR PDF.

While it was shown that the rankings of the lensed candidates were changed by the targeted PDF, these candidates are not yet confirmed as to whether they are true lensed images, and therefore it cannot be said whether the boosting worked as intended.

The changes in rankings were compared to the distribution of the triggers’ $\Delta\phi$ and Δt between Hanford and Livingston detectors, as found from the targeted search run, and their position on the targeted PDF used in the search. These are plotted in figure 15. If these were real lensed image candidates, meaningful conclusions could be drawn between the change in a candidate’s ranking and its position on the PDF. As future work, testing the targeted PDF on a simulated lensed image pair would be able to give results as to whether a lensed sub-threshold image has its ranking boosted by the targeted PDF.

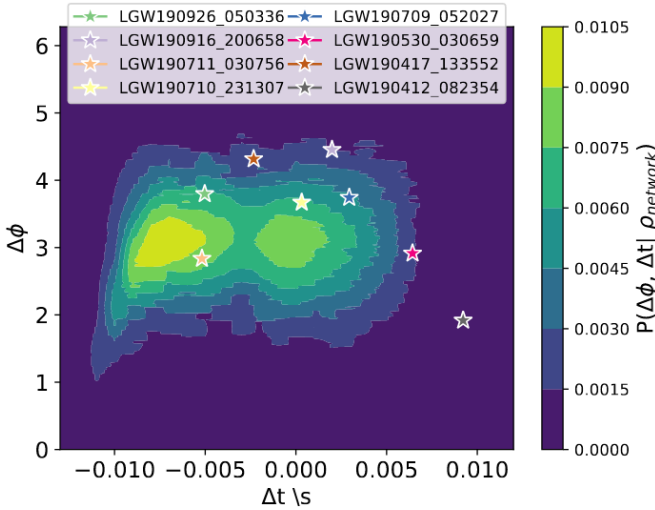


Figure 15. The triggers analysed from the targeted search run as points overlaid onto the targeted $\Delta\phi - \Delta t$ PDF used in the search. The PDF was calculated for detectors Hanford and Livingston, targeted to the 90% credible area of super-threshold event S190408an, and at SNR of 4 in all detectors.

		Whole Sky PDF	Targeted PDF	
S190408an	Likelihood	55.25	35.28	
	FAR	1.189e-25	7.227e-16	
	Ranking	3	4	
LGW190412	Likelihood	16.04	16.7	
	.082354	FAR	5.481e-8	9.089e-8
	Ranking	22	18	
LGW190417	Likelihood	14.23	15.28	
	.133552	FAR	2.884e-7	3.376e-7
	Ranking	31	28	
LGW190530	Likelihood	15.2		
	.030659	FAR	1.225e-7	
	Ranking	23	>100	
LGW190709	Likelihood	14.31	13.86	
	.052027	FAR	2.69e-7	0.00000168
	Ranking	30	51	
LGW190710	Likelihood	16.16	16.72	
	.231307	FAR	5.068e-8	8.928e-8
	Ranking	21	17	
LGW190711	Likelihood	17.13	16.91	
	.030756	FAR	1.568e-8	7.284e-8
	Ranking	18	15	
LGW190916	Likelihood	15.16	16.57	
	.200658	FAR	1.263e-7	9.86e-8
			24	19
LGW190926	Likelihood	14.58	13.8	
	.050336	FAR	2.106e-7	1.736e-6
	Ranking	29	52	

Table I. The original event and eight triggers’ ranking and ranking statistics, log likelihood ratio and FAR, before and after implementing a targeted $\Delta\phi - \Delta t$ PDF using S190408an’s 90% region of its sky localisation.

6. OVERALL CONCLUSIONS

This project aimed to refine one of the current search methods for finding sub-threshold lensed gravitational waves.

Firstly, the waveform family used in the searches was identified, finding that with current uncertainties on spin measurements of CBCs, the effect of choosing a spin-aligned waveform over a precession-included waveform is not expected to significantly alter the detectability of sub-threshold gravitational waves.


Secondly, a constraint was placed on the search pipeline to constrain the search to a particular sky location of a target event, exploiting the large in certainty in gravitational wave separation compared to the expected lensed image separation. The constraint was implemented by the use of a targeted $\Delta\phi - \Delta t$ PDF to the 90% credible area of super-threshold event S190408an, which was used within the search to calculate the log likelihood ratio. A targeted search run was performed on O3a data and the change in rankings of the target event and lensed image candidates were observed. The results from these changes in ranking could not be effectively used to say whether the targeted PDF boosts sub-threshold

lensed images as intended, as the candidates have not yet been evaluated to be true lensed images to the target event. Therefore the refined targeted search needs to be tested on a simulated pair of lensed images, from which it could be seen whether the sky localisation constraint effectively boosts a sub-threshold lensed image of a super-threshold target event.

Once known to be working as expected, this constraint should be able to effectively target lensed candidates with similar sky location to a target event, and derank other triggers which do not meet this criteria. This then would be implemented into the TESLA pipeline and used to redo searches for O1, O2, and O3a results, with a large potential to increase the detectability of sub-threshold lensed gravitational waves.

7. ACKNOWLEDGEMENTS

Thank you so much to Alvin Li and Alan Weinstein for being excellent and encouraging mentors, and both inspiring me with their love of gravitational wave physics. Thank you to Ryan Magee for providing very helpful GstLAL expertise, and all the other amazing faculty and students for creating this amazing summer! Thanks to Caltech, the NSF, and the LIGO lab for helping support this project.



^a storm.colloms@ligo.org

^b kli7@caltech.edu

^c ajw@caltech.edu

- [1] A. Einstein and M. Grossmann, “Covariance properties of the field equations of the theory of gravitation based on the generalized theory of relativity,” *Z. Math. Phys.* **63**, 215–225 (1914).
- [2] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “Properties of the Binary Black Hole Merger GW150914,” *Phys. Rev. Lett.* **116**, 241102 (2016), arXiv:1602.03840 [gr-qc].
- [3] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo during the First and Second Observing Runs,” *Phys. Rev. X* **9**, 031040 (2019), arXiv:1811.12907 [astro-ph.HE].
- [4] R. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “GWTC-2: Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run,” *Phys. Rev. X* **11**, 021053 (2021), arXiv:2010.14527 [gr-qc].
- [5] R. Narayan and M. Bartelmann, “Lectures on gravitational lensing,” *13th Jerusalem Winter School in Theoretical Physics: Formation of Structure in the Universe*, (1996), arXiv:astro-ph/9606001.
- [6] J. Wambsganss, “Gravitational lensing in astronomy,” *Living Reviews in Relativity* **1** (1998), 10.12942/lrr-1998-12.
- [7] D. Kennefick, “Testing relativity from the 1919 eclipse - a question of bias,” *Physics Today - PHYS TODAY* **62** (2009), 10.1063/1.3099578.
- [8] R. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “Search for lensing signatures in the gravitational-wave observations from the first half of LIGO-Virgo’s third observing run,” (2021), arXiv:2105.06384 [gr-qc].
- [9] O. A. Hannuksela, K. Haris, K. K. Y. Ng, S. Kumar, A. K. Mehta, D. Keitel, T. G. F. Li, and P. Ajith, “Search for gravitational lensing signatures in LIGO-Virgo binary black hole events,” *Astrophys. J. Lett.* **874**, L2 (2019), arXiv:1901.02674 [gr-qc].
- [10] C. McIsaac, D. Keitel, T. Collett, I. Harry, S. Mozzon, O. Edy, and D. Bacon, “Search for strongly lensed counterpart images of binary black hole mergers in the first two LIGO observing runs,” *Phys. Rev. D* **102**, 084031 (2020), arXiv:1912.05389 [gr-qc].
- [11] A. K. Y. Li, R. K. L. Lo, S. Sachdev, C. L. Chan, E. T. Lin, T. G. F. Li, and A. J. Weinstein, “Finding diamonds in the rough: Targeted Sub-threshold Search for Strongly-lensed Gravitational-wave Events,” (2019), arXiv:1904.06020 [gr-qc].
- [12] K. Haris, A. K. Mehta, S. Kumar, T. Venumadhav, and P. Ajith, “Identifying strongly lensed gravitational wave signals from binary black hole mergers,” (2018), arXiv:1807.07062 [gr-qc].
- [13] J. M. Ezquiaga, D. E. Holz, W. Hu, M. Lagos, and R. M. Wald, “Phase effects from strong gravitational lensing of gravitational waves,” *Physical Review D* **103** (2021), 10.1103/physrevd.103.064047.
- [14] J. M. Ezquiaga and M. Zumalacárregui, “Gravitational wave lensing beyond general relativity: Birefringence, echoes, and shadows,” *Physical Review D* **102** (2020), 10.1103/physrevd.102.124048.
- [15] K. Kim, J. Lee, R. S. H. Yuen, O. A. Hannuksela, and T. G. F. Li, “Identification of lensed gravitational waves with deep learning,” (2020), arXiv:2010.12093 [gr-qc].
- [16] M. Oguri, “Effect of gravitational lensing on the distribution of gravitational waves from distant binary black hole mergers,” *Mon. Not. Roy. Astron. Soc.* **480**, 3842–3855 (2018), arXiv:1807.02584 [astro-ph.CO].
- [17] S. S. Li, S. Mao, Y. Zhao, and Y. Lu, “Gravitational lensing of gravitational waves: A statistical perspective,” *Mon. Not. Roy. Astron. Soc.* **476**, 2220–2229 (2018), arXiv:1802.05089 [astro-ph.CO].
- [18] K. H. Lai, O. A. Hannuksela, Antonio Herrera-Martín, Jose M. Diego, Tom Broadhurst, and Tjonnie G. F. Li, “Discovering intermediate-mass black hole lenses through gravitational wave lensing,” *Phys. Rev. D* **98**, 083005 (2018), arXiv:1801.07840 [gr-qc].
- [19] R. Takahashi and T. Nakamura, “Wave effects in gravitational lensing of gravitational waves from chirping binaries,” *Astrophys. J.* **595**, 1039–1051 (2003), arXiv:astro-ph/0305055.
- [20] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration, Fermi-GBM, INTEGRAL), “Gravitational Waves and Gamma-rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A,” *Astrophys. J. Lett.* **848**, L13 (2017), arXiv:1710.05834 [astro-ph.HE].
- [21] T. G. F. Li, “Extracting Physics from Gravitational Waves: Testing the Strong-field Dynamics of General Relativity and Inferring the Large-scale Structure of the Universe,” (2013), PhD Thesis, Vrije U., Amsterdam.

- [22] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [23] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “Observation of Gravitational Waves from a Binary Black Hole Merger,” *Phys. Rev. Lett.* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].
- [24] The LIGO Scientific Collaboration and Virgo Collaboration, “Sources and types of gravitational waves,” <https://www.ligo.caltech.edu/page/gw-sources> (2021), accessed on 15.4.2021.
- [25] B. P. Abbott *et al.* (The LIGO Scientific Collaboration and Virgo Collaboration), “GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence,” *Phys. Rev. Lett.* **116**, 241103 (2016), arXiv:1606.04855 [gr-qc].
- [26] S. Sachdev *et al.*, “The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO’s Second and Advanced Virgo’s First Observing Runs,” (2019), arXiv:1901.08580 [gr-qc].
- [27] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D.E. Creighton, R. Everett, *et al.*, “Analysis framework for the prompt discovery of compact binary mergers in gravitational-wave data,” *Physical Review D* **95** (2017), 10.1103/physrevd.95.042001.
- [28] B. C. Lacki, C. S. Kochanek, K. Z. Stanek, N. Inada, and M. Oguri, “Difference Imaging of Lensed Quasar Candidates in the SDSS Supernova Survey Region,” *Astrophys. J.* **698**, 428–438 (2009), arXiv:0801.3432 [astro-ph].
- [29] R. K. L. Lo and I. Magaña Hernandez, “A Bayesian statistical framework for identifying strongly-lensed gravitational-wave signals,” (2021), arXiv:2104.09339 [gr-qc].
- [30] The LIGO Scientific Collaboration and Virgo Collaboration, “Official repositories for tracking gravitational wave event samples,” https://git.ligo.org/pe_event_samples (2021), accessed on 05.07.2021.
- [31] C. Hanna *et al.*, “Fast evaluation of multidetector consistency for real-time gravitational wave searches,” *Phys. Rev. D* **101**, 022003 (2020), arXiv:1901.02227 [gr-qc].
- [32] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, “FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries,” *Phys. Rev. D* **85**, 122006 (2012), arXiv:gr-qc/0509116.
- [33] The LIGO Scientific Collaboration and Virgo Collaboration, “GraceDB G329243 event information,” <https://gracedb.ligo.org/events/G329243/view/>, accessed on 24.09.2021.