Chapter 5

Gravitational Wave Signal Processing

In this chapter, we describe the matched filtering that is done in order to search for a signal of known form within data that also contains noise, as is the case for analyzing GW data.

This problem has been well studied within the field of signal-processing, and the optimal tool for extracting the signal from the noise has been found to be the Weiner Filter [120], also known as the matched or optimal filter.

5.1 The Optimal Filter

Let us assume we have a data stream $s(t)$ that is the sum of a stationary, Gaussian noise $n(t)$ and a signal $h(t)$

$$s(t) = n(t) + h(t) \ . \tag{5.1}$$

Since $n(t)$ is stationary, the mean of $n(t)$ can be taken to zero by defining

$$n(t) \equiv n_{\text{raw}}(t) - \langle n_{\text{raw}} \rangle \ , \tag{5.2}$$

where the brackets denote the expectation value. Since this is true for any stationary series, from here we will assume $n(t)$ has a mean of zero. The two-sided power spectral density (PSD) $S_n(f)$ of this Gaussian noise is defined by

$$\langle \tilde{n}(f)\tilde{n}(f')\rangle \equiv \delta(f-f')S_n(f) \ . \tag{5.3}$$

Let us define a real-valued filter $F(t)$ as

$$A = \int_{-\infty}^{\infty} F(t)a(t)dt = \int_{-\infty}^{\infty} \tilde{F}^*(f)\tilde{a}(f)df \ , \tag{5.4}$$

where A is the filtered value of $a(t)$. Above we have used the convolution theorem and transformed from the time domain to the frequency domain using the Fourier transform of a function, defined by

$$\tilde{a}(f) = \int_{-\infty}^{\infty} a(t)e^{-2\pi i f t}dt \ . \tag{5.5}$$

Using the above properties of the noise, we find

$$\begin{aligned}
\left\langle N^2 \right\rangle &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \tilde{F}^*(f)\tilde{F}^*(f')\left\langle \tilde{n}(f)\tilde{n}(f')\right\rangle df'df \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \tilde{F}^*(f)\tilde{F}^*(f')\delta(f-f')S_n(f)df'df \\
&= \int_{-\infty}^{\infty} \left|\tilde{F}(f)\right|^2 S_n(f)df \ ,
\end{aligned} \tag{5.6}$$

where we have also used the property of real-valued functions that $\tilde{a}^*(f) = \tilde{a}(-f)$ in the last line.

We now wish to find the form of the filter $F(t)$, or in this case $\tilde{F}(f)$, which will optimally extract the signal from the noise. We measure this by maximizing the ratio of filtered values

$$\begin{aligned}
\frac{H^2}{\langle N^2\rangle} &= \frac{\left|\int_{-\infty}^{\infty} \tilde{F}^*(f)\tilde{h}(f)df\right|^2}{\int_{-\infty}^{\infty}\left|\tilde{F}(f)\right|^2 S_n(f)df} \\
&= \frac{\left|\int_{-\infty}^{\infty} \tilde{F}^*(f)\sqrt{S_n(f)}\ \tilde{h}(f)/\sqrt{S_n(f)}df\right|^2}{\int_{-\infty}^{\infty}\left|\tilde{F}(f)\right|^2 S_n(f)df} \ .
\end{aligned} \tag{5.7}$$

We can use the Cauchy-Schwarz inequality,

$$\left| \int_{-\infty}^{\infty} A(f)B(f)df \right|^2 \leq \int_{-\infty}^{\infty} |A(f)|^2 df \int_{-\infty}^{\infty} |B(f)|^2 df \ , \tag{5.8}$$

to argue that in order for the ratio of equation (5.7) to be maximized, the functions corresponding to $A(f)$ and $B(f)$ must be equal up to a constant C, where $A(f) = \tilde{F}^*(f)\sqrt{S_n(f)}$ and $B(f) = \tilde{h}(f)/\sqrt{S_n(f)}$. Equation (5.7) then becomes

$$\begin{aligned} \frac{H^2}{\langle N^2 \rangle} &= \frac{C\left(\int_{-\infty}^{\infty} \left|\tilde{F}(f)\right|^2 S_n(f)df\right)\left(\int_{-\infty}^{\infty} \left|\tilde{h}(f)\right|^2 /S_n(f)df\right)}{\int_{-\infty}^{\infty} \left|\tilde{F}(f)\right|^2 S_n(f)df} \\ &= C\int_{-\infty}^{\infty} \frac{\tilde{h}^*(f)\tilde{h}(f)}{S_n(f)}df \ . \end{aligned} \tag{5.9}$$

From this exercise we find that

$$\tilde{F}^*(f) = C\frac{\tilde{h}(f)}{S_n(f)} \ , \tag{5.10}$$

which is the definition of the optimal filter for $h(t)$ embedded in stationary, Gaussian noise.

Let us use the optimal filter to define an inner product of the data s with the template h:

$$(h|s) = \int_{-\infty}^{\infty} \frac{\tilde{h}^*(f)\tilde{s}(f)}{S_n(f)}df \ . \tag{5.11}$$

An interesting interpretation of equation (5.11) is to split the PSD $S_n(f)$ into two amplitude spectral density (ASD) terms $\sqrt{S_n(f)}$ that can be associated with the data and the template separately,

$$(h|s) = \int_{-\infty}^{\infty} \frac{\tilde{h}^*(f)}{\sqrt{S_n(f)}} \frac{\tilde{s}(f)}{\sqrt{S_n(f)}}df \ . \tag{5.12}$$

This ends up weighting both the template and the data by the inverse of the ASD, which can be seen as "whitening" both the template and the data. The reason this is known as whitening is because for stationary, white, Gaussian noise the PSD is frequency independent resulting in an optimal filter of the form

$$(h|s) = \int_{-\infty}^{\infty} \tilde{h}^*(f)\tilde{s}(f)df \ . \tag{5.13}$$

If we define $\bar{h}(f) = \frac{\tilde{h}^*(f)}{\sqrt{S_n(f)}}$ and $\bar{s}(f) = \frac{\tilde{s}^*(f)}{\sqrt{S_n(f)}}$, equation (5.12) comes into the same form as equation (5.13)

$$(h|s) = \int_{-\infty}^{\infty} \bar{h}^*(f)\bar{s}(f)df \ . \tag{5.14}$$

5.2 The Waveform Overlap

Now that we have an optimal filter for extracting a signal for noisy data, let us use it to define the overlap M of two vectors a and b

$$\mathrm{M} = (a'|b') \ , \tag{5.15}$$

where the "′" denotes a normalization such that

$$a' = \frac{a}{\sigma_a} \ , \tag{5.16a}$$

and

$$\sigma_a^2 = (a|a) \ , \tag{5.16b}$$

and similarly for b.

This definition of the overlap will take its largest value of 1 when a' and b' are the same function, which can be seen by

$$\begin{aligned} (a'|a') &= \left(\frac{a}{\sigma_a}|\frac{a}{\sigma_a}\right) \\ &= \frac{1}{\sigma_a^2}(a|a) \\ &= 1 \ . \end{aligned}$$

Since we are interested in the strength of the signal that matches the template h in the data s, we compute the signal-to-noise ratio (SNR) by normalizing the overlap with respect to the template but not the data

$$\mathrm{SNR} = \frac{1}{\sigma_h} (s|h) \ . \tag{5.17}$$

The value of the SNR will then be proportional to the amplitude of the signal buried in the noise.

So far we have defined the overlap between s and h where we have assumed the template to be of the same length as our data. However, what we are actually interested in is the matched filter of s using h where the data $s(t)$ is an extended time series whose length is longer than the template time series $h(t)$. The matched filter output as a time series is given by

$$(s(t)|h) = \int_{-\infty}^{\infty} \tilde{s}^* \tilde{h} e^{-2\pi i f t} df \ . \tag{5.18}$$

For real-valued time series s and h, this filter outputs a real-valued time series that, in the absence of signal, is χ^2distributed with one degree of freedom.

5.2.1 Matched Filtering for Compact Binary Signals

As derived in chapter 2, the GWs we are searching for come in two polarizations: the plus (+) and cross ($\times$) polarizations. The actual signal seen by a detector is a combination of the two polarizations that can be calculated using the antenna pattern of the detector and the parameters of the source. Specifically, the signal seen by the detector $h(t)$ is given by

$$h(t) = F_+ h_+(t) + F_\times h_\times(t) \ , \tag{5.19}$$

where h_+ and $h_\times$ are the plus and cross polarizations of the signal respectively, and F_+ and $F_\times$ are the antenna patterns of the detector giving the sensitivity to the plus and cross polarizations of GW signals respectively. Descriptions of these are given in chapter 3.

Due to the nature of GW signals from the inspiral phase of CBCs (see chapter 2 for a discussion), we find that the phase evolution of the cross polarization is 90° out of phase with the plus polarization, given by

$$h_+(t) = A(t)\cos\left(\phi(t)\right) \ , \tag{5.20}$$

$$h_\times(t) = A(t)\sin\left(\phi(t)\right) \ , \tag{5.21}$$

where $A(t)$ is the amplitude evolution of the signal and $\phi(t)$ is the phase evolution of the signal. From this, we find a simple relation between the Fourier transform of the two polarizations:

$$\tilde{h}_+ = i\tilde{h}_\times \ . \tag{5.22}$$

When we filter $s = (Xh_+/\sigma_h) + (Yh_\times/\sigma_h)$ with the template h_+, we obtain the matched-filter real-time series z_+

$$z_+ = (s|h_+) = \frac{X}{\sigma_h}\left(h_+|h_+\right) + \frac{Y}{\sigma_h}\left(h_\times|h_+\right) \ . \tag{5.23}$$

Using equation (5.22), we find

$$\begin{aligned}
(h_\times|h_+) &= \int_{-\infty}^{\infty} \frac{\tilde{h}_\times^* h_+ + \tilde{h}_+^* h_\times}{S_n} df \\
&= \int_{-\infty}^{\infty} \frac{\left(i\tilde{h}_+\right)^* h_+ + \tilde{h}_+^* i h_+}{S_n} df \\
&= \int_{-\infty}^{\infty} \frac{-i\tilde{h}_+^* h_+ + \tilde{h}_+^* i h_+}{S_n} df \\
&= 0 \ ,
\end{aligned}$$

which then implies

$$z_+ = X\sigma_h \ , \tag{5.24a}$$

where we have used equation (5.16b) for the definition of σ_h. A similar procedure shows that when we filter s with the template $h_\times$, we obtain the matched filter time series $z_\times$

$$z_\times = Y\sigma_h \ . \tag{5.24b}$$

From equation (5.24) we construct our SNR ρ using the combination

$$\rho = \frac{1}{\sigma_h}\sqrt{|z_+|^2 + |z_\times|^2}\ . \tag{5.25}$$

This is referred to as a *two-phase filter*, which has twice the degrees of freedom of a single-phase filter. The bonus of extracting information from both polarizations of the GW comes with the cost of increasing the expectation value of ρ when there is no signal present.

By combining the definition of the SNR (equation (5.25)) with the template normalization σ_h (equation (5.16b)) we can define an effective distance for a given trigger

$$D_{\text{eff}} = \frac{\sigma_h}{\rho}\ . \tag{5.26}$$

We choose a normalization distance of 1 Mpc to set the templates' amplitudes. This results in the units of the effective distance being Mpc.

5.2.2 Template Bank Construction

The manifold of waveforms is a continuous space in the component masses, of which we are only able to search discrete points. In order to make sure we do not miss a signal because its parameters are slightly different from what we are searching for, we construct a bank of templates in such a way as to minimize the loss of a signal's SNR. This is done by computing the overlap of waveforms with different parameters using equation (5.15).

As noted above, the overlap of a waveform with itself is unity. Any *mismatch*, $1 - \mathrm{M}$, between two waveforms will show up as a reduction of the recovered SNR when searching for one waveform with the other. When constructing a template bank, it is useful to view this mismatch as a measure of the distance between two templates. With this in mind, we can create a metric that will tell us the distance between two templates of different parameters. This is done by defining the parameter space in which the metric exists, and then computing the mismatch between templates. For two templates infinitesimally far apart in the parameter space θ^μ, the mismatch is [121, 122]

$$\begin{aligned}
1 - \mathrm{M} &\approx 1 - \left(h\left(\theta^\mu + d\theta^\mu\right)|h\left(\theta^\nu\right)\right) \\
&\approx \frac{1}{2}\left[1 + 1 - 2\left(h\left(\theta^\mu + d\theta^\mu\right)|h\left(\theta^\nu\right)\right)\right] \\
&\approx \frac{1}{2}[\left(h\left(\theta^\mu\right)|h\left(\theta^\nu\right)\right) + \left(h\left(\theta^\mu + d\theta^\mu\right)|h\left(\theta^\nu + d\theta^\nu\right)\right) \qquad (5.27)\\
&\quad - \left(h\left(\theta^\mu + d\theta^\mu\right)|h\left(\theta^\nu\right)\right) - \left(h\left(\theta^\mu\right)|h\left(\theta^\nu + d\theta^\nu\right)\right)] \\
&\approx \frac{1}{2}\left[\left(h\left(\theta^\mu + d\theta^\mu\right) - h\left(\theta^\mu\right)|h\left(\theta^\nu + d\theta^\nu\right) - h\left(\theta^\nu\right)\right)\right] \\
&\approx \frac{1}{2}\left[\left(\frac{\partial h}{\partial \theta^\mu} d\theta^\mu | \frac{\partial h}{\partial \theta^\nu} d\theta^\nu\right)\right] \\
&\approx \frac{1}{2}\left[\left(h_\mu|h_\nu\right)\right] d\theta^\mu d\theta^\nu \\
&\approx \frac{1}{2} g_{\mu\nu} d\theta^\mu d\theta^\nu \ , \qquad (5.28)
\end{aligned}$$

where $h_\mu \equiv \frac{\partial h}{\partial \theta^\mu}$, and we have used the symmetry $(a|b) = (b|a)$. The above equation defines the metric $g_{\mu\nu}$ to be of the quadratic form

$$g_{\mu\nu} \equiv \left(h_\mu|h_\nu\right) \ . \qquad (5.29)$$

The metric (equation (5.29)) can be used to predict to quadratic order how quickly the mismatch grows as we move in different directions in parameter space away from a particular point. Using this metric we place templates such that the furthest distance any point is from a template is less than a tolerance value ϵ. This will ensure that no more than ϵ of the SNR is lost due to our discretization of parameter space.

The question we must now ask is which parameter space to use. Answering this question specifies which directions we will use when taking partial derivatives of the templates. In the end, it is useful to choose those directions such that the metric is as flat as possible across the space. This helps in developing an algorithm that most efficiently covers the parameter space. Such a space has been found (equations (23) and (24) in Ref. [123]) whose directions correspond to the *chirp times* of a waveform. Using 2.0 PN order waveforms for nonspinning binary objects in quasi-circular orbital evolution, there are chirp times τ_i for $i = 0, 2, 3, 4$, corresponding to the chirp times proportional to $(v/c))^i$ in the PN expansion of the phase. These chirp times are given in terms of the mass parameters as [124, 125, 122]

$$\tau_0 = \frac{5}{256\,(\pi f_0)^{8/3}}\mathcal{M}^{-5/3}\,, \tag{5.30a}$$

$$\tau_2 = \frac{5}{192\eta^{2/5}\,(\pi f_0)^2}\left(\frac{743}{336} + \frac{11}{4}\eta\right)\mathcal{M}^{-1}\,, \tag{5.30b}$$

$$\tau_3 = \frac{\pi}{8\eta^{3/5}\,(\pi f_0)^{5/3}}\mathcal{M}^{-2/3}\,, \tag{5.30c}$$

$$\tau_4 = \frac{5}{128\eta^{4/5}\,(\pi f_0)^{4/3}}\left(\frac{3\,058\,673}{1\,016\,064} + \frac{5429}{1008}\eta + \frac{617}{144}\eta^2\right)\mathcal{M}^{-1/3}\,, \tag{5.30d}$$

where M is the total mass of the binary, $\eta = m_1 m_2/M^2$ is the symmetric mass ratio, $\mathcal{M} = \eta^{3/5}M$ is the chirp mass, and f_0 is some fiducial frequency. If f_0 is chosen to be the starting frequency for the waveform generation, then the chirp times correspond to the 0PN length of the chirp for τ_0, and the $(i/2)$PN correction to the chirp time for $i = 2, 3, 4$. For the search described in later chapters, we have chosen to use $(\tau_0,\ \tau_3)$ as the parameter space in which we lay our templates, however we compute a 3D metric that additionally includes the difference in coalescence times of two waveforms for use in other portions of the search. Although the metric is not perfectly flat for PN order greater than 1.0PN, we assume it is slowly varying, and thus flat in the local vicinity of a point. When constructing the bank, we use the hexagonal placement algorithm [123], which is the most efficient placement algorithm for a flat, 2D space.

5.2.3 Signal-Based Vetoes

Signal-based vetoes are discriminators that can be used to separate triggers arising from either a random instantiation of the Gaussian noise or transient excess power in the data from triggers arising from actual signals. These vetoes work by comparing how a particular aspect of the data should behave in the presence of a signal and not.

In the following sections we discuss several signal-based vetoes that look at different aspects of the data. The first two signal-based vetoes described below, the $\chi^2$and r^2vetoes, use information from a single detector. Since we require triggers to be coincident between multiple detectors, which we will discuss in more detail in chapter 6, there are addition signal-based vetoes we can perform using information from more than one detector. The effective distance cut and the amplitude consistency check are examples of this, using information from two coaligned detectors.

5.2.3.1 The χ^2Veto Calculation

The χ^2veto tests the consistency of the data with what we expect from a signal by looking at how the SNR, ρ, is accumulated from different parts a template. This is done by breaking up the template into p continuous bins such that each is expected to provide ρ/p to the SNR calculation in the presence of a real signal. This calculation is described by

$$\chi^2 = \sum_{i=1}^{p} \left(\rho_i - \frac{\rho}{p} \right)^2 , \tag{5.31}$$

where p is the number of χ^2bins, ρ_i is the SNR from the ith bin of the template, and ρ is the total SNR.

With this formulation, in the presence of Gaussian noise, $\chi^2$is a quantity that is χ^2distributed with $2p-2$ degrees of freedom:

$$\left\langle \chi^2 \right\rangle_{\mathrm{N}} = 2p - 2 \; . \tag{5.32}$$

Normally such a quantity would be χ^2distributed with $p-1$ degrees of freedom, however there is an extra factor of 2 because ρ^2 is calculated from a two-phase filter, which has an expectation value of $\left\langle \rho^2 \right\rangle = 2$ in Gaussian noise.

The $\chi^2$is not actually the quantity we threshold on when applying the χ^2veto. This is because the expectation value of χ^2in the presence of signal is not the same as in Gaussian noise. Since we are using a discrete template bank to search for signals in a continuous parameter space, we will not recover a signal with a template of those exact parameters. If we calculate the expectation value of χ^2in the presence of signal without noise, evaluated at the time when the SNR is maximized, we find

$$\langle \chi^2 \rangle_{\mathrm{S}} = \delta\rho^2 \ , \tag{5.33}$$

where δ is the mismatch between the template and the signal. Since a real signal can have a large value of $\chi^2$for high SNRs, we end up thresholding on a normalized χ^2, ξ^2, where we veto triggers with

$$\xi^2 \equiv \frac{\chi^2}{p + \delta\rho^2} > \xi_*^2 \ . \tag{5.34}$$

For the χ^2veto, p, δ, and ξ_* are tunable parameters.

It should be noted that in order to do this calculation, p additional inverse Fourier Transforms of the data need to be computed for the p different parts of the template.

5.2.3.2 The r^2Veto Calculation

Another signal-based veto that has been developed is the r^2veto. This veto is also a measure of how much the data looks like an actual signal. The difference between the r^2veto and the χ^2veto is that the r^2veto looks over a stretch of data rather than at a single point in time. Specifically, the r^2veto measures the amount of time the χ^2time series spends above a particular threshold for the T seconds prior to an trigger. The duration r^2_{duration} is given as

$$r^2_{\mathrm{duration}} = \frac{d}{f_{\mathrm{sampling}}} \ , \tag{5.35}$$

where d is the number of data points in the time window $(t_0 - T, T)$ that have $\chi^2/p > r_*^2$, and f_{sampling} is the sampling rate of the data. For the r^2veto, T and r_*^2 are tunable parameters.

5.2.3.3 The Effective Distance Cut

When looking for coincident triggers between detectors that are coaligned, we apply an effective distance cut. Coaligned detectors, such as the two collocated Hanford detectors H1 and H2, have the interesting property of sharing the same antenna patterns. The effective distance seen by a particular detector for a given CBC GW signal is

$$D_{\text{eff}} = \frac{D}{\sqrt{F_+^2 \left(1+\cos^2 \iota\right)^2/4 + F_\times^2 \cos^2 \iota}}\,, \tag{5.36}$$

which is only dependent on the distance to the source D, the inclination angle ι between the orbital plane and the line of site between the source and the detector, and the detector antenna patterns F_+ and $F_\times$. In terms of measured quantities, the effective distance is defined in equation (5.26). Since coaligned detectors have the same antenna patterns, they should see the same D_{eff} for triggers coming from a real CBC GW signal.

Since noise fluctuations in the different detectors can separately change the effective distance seen by each detector, we do not require the effective distances to agree perfectly. Instead we allow for some variation and veto coincident triggers that have a fractional effective distance difference κ greater than a given threshold κ_* (i.e., $\kappa > \kappa_*$). We compute κ as

$$\kappa = \frac{2\left|D_{\text{eff,A}} - D_{\text{eff,B}}\right|}{D_{\text{eff,A}} + D_{\text{eff,B}}}\,, \tag{5.37}$$

where $D_{\text{eff,A}}$ is the effective distance from detector A, $D_{\text{eff,B}}$ is the effective distance from detector B.

This veto is only applied between coaligned detectors due to the fact that the same signal can be seen by two nonaligned detectors with very different amplitudes, and thus effective distances, since the blind spots of the detectors' antenna patterns do not coincide. The blind spots are a very small portion of the sky and thus it is not unreasonable to apply a similar veto between nonaligned detectors. In that case, the fraction of true signals being vetoed can be calculated by looking at the volume of the universe in which such a signal would originate, bounded by the distance at which each detector would see a signal with an SNR above a particular threshold.

5.2.3.4 The Amplitude Consistency Check

In this section we describe the calculation of an amplitude consistency check. This check is very similar to the effective distance cut, however instead of checking the consistency between two triggers that make up a coincident trigger, we are checking whether the lack of a trigger from one coaligned detector is consistent with the measured parameters of the trigger in the other detector.

As always, we use the PSD from each interferometer to calculate a template normalization σ (equation (5.16b)) for a given time. With this, we compute a horizon distance for a canonical mass BNS template as the distance to which we would see an optimally located and oriented signal with an SNR of 8. Similar to the effective distance (equation (5.26)), this is given as

$$D_{\mathrm{horizon}} = \frac{\sigma}{8} \; . \tag{5.38}$$

We can rearrange the effective distance cut (equation (5.37)) using the effective distance (equation (5.26)) for the trigger from detector A and the horizon distance (equation (5.38)) from detector B. Solving for ρ_{A}, we find the veto to be

$$\rho_{\mathrm{A}} > \left(\frac{2+\kappa_*}{2-\kappa_*}\right)\left(\frac{D_{\mathrm{horizon,A}}}{D_{\mathrm{horizon,B}}}\right)\rho_{\mathrm{threshold,B}} \; . \tag{5.39}$$

In order to use this equation to veto triggers from detector A when there is no coincident trigger in detector B, we have substituted the SNR threshold of detector B $\rho_{\mathrm{threshold,B}}$ in the place of ρ_{B}. This veto tells us to veto triggers from detector A when the ratio of the horizon distances in detectors A and B is small enough such that there should have also been a trigger above the SNR threshold of detector B. The factor $(2+\kappa_*)\,/\,(2-\kappa_*)$ allows for noise fluctuations to change ρ_{A}, $D_{\mathrm{horizon,A}}$, and $D_{\mathrm{horizon,B}}$ by some amount.

This veto is also only applied between coaligned detectors due to the fact that the there exist locations in the sky and polarizations of a signal such that for two different, nonaligned detectors, one of them should see the signal above the SNR threshold, while the other should not. As in the case for the effective distance cut, a similar veto can be constructed for nonaligned detectors if one is willing to veto a fraction of true signals.

5.3 Previous Detection Statistic: Effective SNR

For the previous searches, a detection statistic was developed to separate signals from background. This statistic is the combined effective SNR and is constructed as follows.

The single-detector SNR ρ is produced by matched filtering the data against our templates. From ρ we define the *effective SNR*, ρ_{eff}, as

$$\rho_{\mathrm{eff}}^2 = \frac{\rho^2}{\sqrt{\left(\frac{\chi^2}{2p-2}\right)\left(1 + \frac{\rho^2}{250}\right)}} \,, \tag{5.40}$$

where p is the number of bins used by the χ^2veto, $2p-2$ is the number of degrees of freedom of the χ^2veto in Gaussian noise, and the 250 is a tunable parameter that helps to further separate signals from background. This definition of the effective SNR reduces the apparent significance of nonGaussian instrumental artifacts since it weights the SNR by the χ^2. This effectively reduces the significance of outliers to the expected SNR distribution due to Gaussian noise (i.e., nonGaussian instrument artifacts) while minimally affecting the apparent significance triggers of real signals.

We then combine the effective SNRs for the single-detector triggers that form a coincident trigger into the *combined effective SNR*, ρ_{c}, for that coincident trigger using

$$\rho_{\mathrm{c}}^2 = \sum_{i=1}^{N} \rho_{\mathrm{eff},i}^2 \;. \tag{5.41}$$