

This is the accepted manuscript made available via CHORUS. The article has been published as:

Detecting gravitational waves in data with non-stationary and non-Gaussian noise

Barak Zackay, Tejaswi Venumadhav, Javier Roulet, Liang Dai, and Matias Zaldarriaga

Phys. Rev. D **104**, 063034 — Published 20 September 2021

DOI: [10.1103/PhysRevD.104.063034](https://doi.org/10.1103/PhysRevD.104.063034)

Detecting Gravitational Waves in Data with Non-Stationary and Non-Gaussian Noise

Barak Zackay,^{1,*} Tejaswi Venumadhav,¹ Javier Roulet,² Liang Dai,¹ and Matias Zaldarriaga¹

¹*School of Natural Sciences, Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA*

²*Department of Physics, Princeton University, Princeton, NJ, 08540, USA*

(Dated: July 26, 2021)

Searches for gravitational waves crucially depend on exact signal processing of noisy strain data from gravitational wave detectors, which are known to exhibit significant non-stationary and non-Gaussian behavior. In this paper, we study two distinct effects in the LIGO/Virgo data which reduce the sensitivity of searches: first, variations in the noise power spectral density (PSD) on timescales of more than a few seconds; and second, loud and abrupt transient ‘glitches’ of terrestrial or instrumental origin. We derive a simple procedure to correct, at first order, the effect of the variation in the PSD on the search background. Given the knowledge of the existence of localized glitches in particular segments of data, we also develop a method to insulate statistical inference from these glitches, so as to cleanly excise them without affecting the search background in neighboring seconds. We show the importance of applying these methods on the publicly available LIGO data, and estimate an increase in the detection volume of at least 15% from the PSD-drift correction alone, due to the improved background distribution.

I. INTRODUCTION

The first detection of Gravitational Waves (GWs) from merging compact binaries in LIGO data opened up the opportunity to study a new family of astrophysical sources. Multiple sources have been detected in the first and second observing runs of Advanced LIGO and Virgo (O1 and O2) [1–9]. State-of-the-art searches [10–13] suffer from reduced sensitivity due to non-Gaussian effects in the strain data [14], which results in fewer detected sources. It is of paramount importance to understand and correct for these non-Gaussian effects in order to maximize the yield of the existing data.

We distinguish two distinct effects that contribute to the tail distribution of triggers obtained by matched filtering with a template bank. The first, non-stationarity, is due to the changing power spectral density (PSD) of the noise that constitutes the background for the GW search. As we show, changes in the PSD on short timescales (even as short as 10s) are important to capture in order to suppress over-production of candidate signals with high signal-to-noise ratio (SNR). The second effect, non-Gaussian behavior, consists of abrupt noise transients (‘glitches’ on sub-second timescales) that are caused by either environmental disturbance or instrumental malfunction. The origin of many of these glitches is not yet understood [14].

In this paper, we treat the problem of detecting putative GW events in the presence of these systematic effects. Firstly, we propose a simple practical solution that corrects the effect of a varying PSD on short timescales, to first order in the change in the PSD. Secondly, we present a method to isolate the detection statistic for compact binaries from abrupt noise transients at given times and of given lengths.

Section II is devoted to the effect of slow variations in the PSD. We begin by reviewing matched filtering, and the process of PSD estimation, in stationary Gaussian noise. We demonstrate the magnitude and timescale of the PSD variations in the LIGO data, and derive how these variations cause a loss in sensitivity for searches. We then propose a practical and simple way to cancel the first order loss in sensitivity due to the mis-estimation of the local PSD. Essentially, this solution is to divide all computed matched filtering overlaps by their locally estimated standard deviation $\sigma_z(t)$, given by:

$$\sigma_z^2(t) = \frac{1}{N_a} \sum_{t'=t-N_a\Delta t/2}^{t+N_a\Delta t/2} |z(t')|^2, \quad (1)$$

where $z(t)$ are the matched filtering overlaps, defined in Section II A, and N_a is the number of scores used to estimate the average. Finally, we use GW triggers on the entirety of the publicly available O2 LIGO data [15, 16] to explicitly demonstrate the dramatic effects of PSD changes on the tail distribution of matched filtering scores, and quantify the sensitivity gain when we apply this simple correction to the overlaps.

In Section III we explain how to null the effect of identified abrupt noise transients on matched filtering scores of templates for long GW signals. We achieve this by replacing (or ‘inpainting’) the bad segments of strain data with values (that we solve for), such that the inverse-PSD-filtered (i.e., twice-whitened, or ‘blued’) data is zero at all bad times. This guarantees that the offending data has zero influence on any computation of the likelihood, and preserves optimal sensitivity for any real GW event in the surrounding good data.

The methods described in this paper have been implemented in the search pipeline described in Ref. [13], together with other improvements in candidate ranking [9], signal consistency checks [17] and template bank subdivision [18].

* bzackay@ias.edu

II. VARIATIONS IN THE NOISE PSD

In this section we demonstrate that the noise PSD of the LIGO data exhibits slow drifts over timescales of more than a few seconds, and discuss the loss in sensitivity due to this drift. We present a simple way to mitigate this problem and quantify the gain in search volume from the proposed correction.

A. PSD estimation and matched filtering

We begin by defining and briefly reviewing the essential features of matched filtering for signals within data with stationary Gaussian random noise. Similar derivations in past literature could be found in [19–21]. The statistical properties of the noise are completely described by its autocorrelation function, $C_n(\tau)$, defined by

$$\langle d(t) d(t + \tau) \rangle = C_n(\tau). \quad (2)$$

The stationary nature manifests in the fact that the function C_n is only a function of the lag τ . Additionally, $C_n(\tau)$ decays to zero for large values of the lag. The two-sided Power Spectral Density (PSD) is defined as the continuous Fourier transform of the autocorrelation function, i.e.,

$$S_{n,2}(f) = \int_{-\infty}^{\infty} d\tau C_n(\tau) e^{-i2\pi f\tau}. \quad (3)$$

The data is real-valued, and hence $C_n(\tau)$ and $S_{n,2}(f)$ are real-valued and even functions of their arguments. It is conventional to define the one-sided PSD as $S_n(f) = S_{n,2}(f) + S_{n,2}(-f) = 2 S_{n,2}(f)$.

We are interested in searching for a signal, say $h(t)$, within the data. In our use-case, we can take $h(t)$ to have compact support when restricted to the sensitive band of the detectors. We work with data that is sampled at an interval Δt , chosen such that the Nyquist frequency $f_s/2 = 1/(2\Delta t)$ is high enough to encompass the sensitive band.

Let us consider a segment of data, of length N , that is long enough to contain the putative signal. In the absence of the signal, the Discrete Fourier Transform (DFT) of the segment satisfies¹

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = \frac{1}{2} e^{i\pi(f_{m'} - f_m)\Delta t}$$

¹ Our convention for the DFT is

$$\tilde{d}(f_m) = \sum_{n=0}^N d(n\Delta t) e^{-2\pi i n \Delta t f_m},$$

where

$$f_m = \frac{m}{N\Delta t}, \quad -N/2 + 1 \leq m \leq N/2,$$

with N being even.

$$\times \int_{-\infty}^{\infty} df S_{n,2}(f) W_N(f; f_m) W_N(f; f_{m'}), \quad (4)$$

where the window function $W_N(f; f_m)$ is defined in Eq. (A6) (see Appendix A1 for a derivation). When viewed as a function of frequency f , $W_N(f; f_m)$ exhibits a series of peaks with height N and width $\sim 1/(N\Delta t)$, separated by the sampling frequency. If we assume that the data was properly bandpassed before sampling (to prevent aliasing), we can restrict to the frequency interval between $-f_s/2$ and $f_s/2$. If the PSD behaves smoothly on frequency scales of $\sim 1/(N\Delta t)$, $W_N(f; f_m)$ behaves like a delta-function selecting the frequency $f = -f_m$ in the integrand, and we have

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle \approx \frac{N}{2\Delta t} S_n(f_m) \delta_{m,m'}. \quad (5)$$

It is convenient to define a whitened data stream, d_w , according to:

$$\widetilde{d_w}(f_m) = \left[\frac{2\Delta t}{S_n(f_m)} \right]^{1/2} \tilde{d}(f_m), \quad (6)$$

that satisfies

$$\left\langle \widetilde{d_w}(f_m) \left[\widetilde{d_w}(f_{m'}) \right]^* \right\rangle = N \delta_{m,m'}. \quad (7)$$

The whitened data is the result of convolving the raw data stream with a whitening filter. We define the convolution of the signal and the whitening filter as the ‘whitened signal’, h_w :

$$\widetilde{h_w}(f_m) = \left[\frac{2\Delta t}{S_n(f_m)} \right]^{1/2} \tilde{h}(f_m). \quad (8)$$

Under the assumption of stationary Gaussian noise, the matched filtering score, defined as

$$z = \left[d_w \circledast \overleftarrow{h_w} \right] (0) = \frac{1}{N} \sum_m \left[\widetilde{h_w}(f_m) \right]^* \widetilde{d_w}(f_m) \quad (9)$$

is the optimal detection statistic for the signal $h(t)$ in the data (in the first equation, the symbol \circledast represents convolution and the arrow signifies time-reversal). If the data d and signal h are real-valued, so is the score z . In the absence of a signal, the variance of the score is

$$\langle z^2 \rangle = \left[h_w \circledast \overleftarrow{h_w} \right] (0) = \frac{1}{N} \sum_m \left| \widetilde{h_w}(f_m) \right|^2, \quad (10)$$

which also equals its expectation value in the presence of the signal. Hence the quantity in Eq. (10) is the squared signal-to-noise ratio (SNR^2). We can search for signals with different arrival times by enumerating over shifts in $h(t)$ in Eq. (9) using an implementation of the Fast Fourier Transform (FFT). In practice, we also search for signals with constant phase-offsets, in which case the template $h(t)$ is complex-valued, and the detection statistic is the absolute value of Eq. (9). For the sake of simplicity, we assume that the template is real; all the results of this paper hold for complex templates as well.

In deriving Equation (9) and in proving its optimality, it is assumed that the noise PSD, $S_n(f)$, is known, while in practice, we have to measure it from the data itself. A standard way to do so is the Welch method, which divides the data into many (ideally overlapping) segments, applies a window function followed by a DFT to each segment, and calculates $S_n(f)$ as the average of the power spectra in all the segments.

The frequency resolution and the precision of the measured PSD are important criteria to decide the number and duration of segments to use with the Welch method. The frequency resolution depends on the length of the segments, through the window functions $W_N(f; f_m)$ defined in Eq. (A6). The advanced LIGO noise contains several sharp spectral lines, at which the PSD is several orders of magnitude higher than the ‘floor’. If the chosen segments are too short, the lines are broadened (through convolution with $W_N(f; f_m)$) and bleed into surrounding frequency bins: the effect is to reduce the SNR, and move us away from optimality. The stochastic error of the PSD measurement depends on the number of segments that are averaged over. As shown in the next section, the signal recovery efficiency in the presence of stochastic PSD errors is roughly $1 - 0.5 N_{\text{seg}}^{-1}$, where N_{seg} is the number of segments used when estimating a PSD using the Welch method. Together, these parameters define a minimal duration over which the PSD needs to be estimated to achieve a target recovery efficiency. In order to bound the sensitivity loss to less than a few percent, for LIGO data, the duration of the PSD measurement needs to be roughly 10^3 s.

The discussion so far has been theoretical; we would like to test our assumptions, and whether the above procedure achieves the required bounds on the sensitivity. The most direct check is to verify the statistics of the matched filtering scores, since they determine our sensitivity. In what follows, we normalize the template h so that the variance $\langle z^2 \rangle$ as given by Eq. (10) equals unity.

We consider the matched filtering scores for a given template $h(t)$ as a time-series, i.e.,

$$\begin{aligned} z(t = n\Delta t) &= \left[d_w \otimes \overleftarrow{h_w} \right] (t) \\ &= \frac{1}{N} \sum_m \left[\widetilde{h_w}(f_m) \right]^* \widetilde{d_w}(f_m) e^{2\pi i f_m n \Delta t}. \end{aligned} \quad (11)$$

The simplest check is whether the actual variance of the scores (as estimated from the time-series) is consistent with unity. Figure 1 shows the histogram of the estimated variance of the scores z for a single template over all the valid² segments of O2 data. We estimated the variance by averaging the power z^2 within rolling windows of length ~ 15 s, which should achieve 2% stochastic error




FIG. 1: Histogram of variance values of overlaps computed on all the valid O2 data. Measurement error on these values is 2%. Measured deviations from unity are much bigger (about 10%) and therefore produce a big difference in significance determination if unaccounted for.

on the variance. Instead, we see that the standard deviation of the variance distribution is approximately 8–9%. In the next section, we investigate the associated loss in sensitivity in more detail. We now investigate the cause of the phenomenon shown in Fig. 1.

The derivation of the variance in Eq. (10) depends on the noise being described by the PSD $S_n(f)$, i.e., the whitened data satisfying Eq. (7). The failure in Fig. 1 suggests that the PSD we estimated did not whiten the data perfectly. It is well known that the behavior of the detector varies with time (a drastic example of this is the scaling of the noise curve with the varying level of human activity in the vicinity of the detector). Due to these phenomena, the noise characteristics can change over timescales that are shorter than the $\mathcal{O}(10^3)$ seconds we use to measure the PSD.

We can view the variance of this series (or instantaneous power, z^2) itself as a time-series, that is described by its own PSD, S_{z^2} . In the stationary case (when Eq. (5) holds), S_{z^2} equals

$$S_{z^2}(f_m \neq 0) = 4\Delta t \left| \widetilde{h_w \otimes \overleftarrow{h_w}} \right|^2(f_m), \quad (12)$$

where $h_w \otimes \overleftarrow{h_w}$ is the autocorrelation function of the whitened waveform (see Appendix A 2 for a derivation). The autocorrelation function of a typical waveform has a width that is of order few ms, and hence at frequencies smaller than 1 Hz we expect the PSD of the power, S_{z^2} , to be flat.

Figure 2 shows the PSD of the variance S_{z^2} of the matched filtering scores, computed using a heavy binary black hole template (we estimate the local variance by convolving the z^2 series with a rolling rectangular window of 1 s duration). The dashed curves are the estimated S_{z^2} , for scores measured on stationary Gaussian noise

² Overlaps in places where our pipeline [13] found local excess noise were not included in Fig 1.




FIG. 2: It is necessary to track the drifting PSD on time scales of seconds. Solid lines show the empirically measured power spectrum S_{z2} of a time-series composed of the measured variance of the overlaps in every second (note that the frequencies are much lower than the frequency content of the waveform itself). For reference, dashed lines show the same on artificially generated stationary Gaussian noise. The variance time-series has a red-noise power spectrum.

generated using the fiducial PSDs for the Livingston and Hanford detectors, respectively: they are flat as a function of frequency, in line with the prediction of Eq. (12). The solid curves show the empirically measured variance PSDs, S_{z2} , for L1 and H1 data, averaged over the valid of O2. We omitted from the average any region flagged as invalid either by the LIGO and Virgo Collaboration, or by our pipeline (see [13]), keeping only the contiguous segments. In the rest of this section, we describe how non-stationary noise can lead to such red-noise spectra, and what the measured curve tells us about the departure from the stationary case.

In the non-stationary case, we start by generalizing Eqs. (2) and (3):

$$\langle d(t - \tau/2) d(t + \tau/2) \rangle = C_n(\tau) + \delta C_n(\tau; t), \quad (13)$$

and

$$S_{n,2}(f; t) = S_{n,2}(f) [1 + \epsilon(f; t)], \quad \text{where} \quad (14)$$

$$\epsilon(f; t) = \frac{1}{S_{n,2}(f)} \int_{-\infty}^{\infty} d\tau \delta C_n(\tau; t) e^{-i2\pi f\tau}. \quad (15)$$

In the above equations, $\epsilon(f; t)$ is the fractional change in the noise PSD at frequency f , which we take to vary on timescales much larger than τ (the frequency f is conjugate to the short timescale τ).

If we whiten the data using the PSD estimated assuming stationary noise (i.e., $S_n(f)$) over a duration that is longer than the timescales over which the PSD varies, the equivalent of Eq. (7) is:

$$\langle \widetilde{d}_w(f_m) [\widetilde{d}_w(f_{m'})]^* \rangle \approx N \delta_{m,m'} + \tilde{\epsilon}(\bar{f}; \Delta f), \quad (16)$$

where $\bar{f} = (f_m + f_{m'})/2$, and $\Delta f = f_m - f_{m'}$ (see Appendix A 1 for a derivation). In the second term, $\tilde{\epsilon}(\bar{f}; \Delta f)$ is the DFT, evaluated at Δf , of $\epsilon(\bar{f}, T = n\Delta t)$ sampled at a rate of $f_s = 1/(\Delta t)$. Equation (16) was derived under the assumption that the noise PSD, $S_n(f)$, is smooth on frequency scales of Δf , and as such is inaccurate in the immediate vicinity of spectral lines in $S_n(f)$.

The non-stationary term, ϵ , correlates Fourier modes of the data, $d(f_m)$, with different frequencies. Intuitively, if we analyze segments that are shorter than the slow timescale, T , and whiten them using the instantaneous PSD, correlations between Fourier modes are diagonal in terms of frequency f_m (as in Eq. (7)). Over longer timescales, the frequencies of the (slow) PSD variations beat against the frequencies of the (fast) Fourier modes and lead to the second term in Eq. (16).

We are interested in the effect of the non-stationary part of the noise PSD, ϵ , on the PSD of the power in the matched filtering overlaps, S_{z2} . For this, it is useful to view the values of $\epsilon(f_m, T)$ themselves as being drawn from a set of random time-series (one for each ‘fast’ frequency f_m). The simplest model for this series is that all the $\epsilon(f_m, T)$ vary in step with each other, and with the same amplitude (it is a straightforward generalization to model more complicated behavior). In this case, there is a single PSD that describes the variations:

$$\langle \tilde{\epsilon}(f_m, f_a) [\tilde{\epsilon}(f_{m'}, f_b)]^* \rangle = \frac{N}{2\Delta t} S_\epsilon(f_a) \delta_{a,b}. \quad (17)$$

We can obtain a simple form for the corrected version of Eq. (12) under additional approximations: (a) variations in the noise PSD (described by S_ϵ) have support at much lower frequencies than the whitened signal \widetilde{h}_w does, and (b) the time-domain whitened waveform is much shorter than the timescales over which $\epsilon(f; T)$ varies. In this case, we have

$$\begin{aligned} S_{z2}(f_m \neq 0) & \approx 2\Delta t \left| \widetilde{h}_w \otimes \widetilde{h}_w \right|^2(f_m) \left[2 + \int_{-\infty}^{\infty} df S_\epsilon(f) \right] + \\ & \left| \widetilde{h}_w^2(f_m) \right|^2 S_\epsilon(f_m). \end{aligned} \quad (18)$$

Appendix A 2 presents a detailed derivation of this equation. The first term in Eq. (18) is flat with frequency at low frequencies (similarly as in Eq. (12)). The second term is proportional to the power spectrum, S_ϵ , of the non-stationary part of the noise-PSD, ϵ . The behavior of the solid curves in Fig. 2 suggests that the PSD, $S_n(f; T)$, itself varies on timescales T larger than a few seconds, and that these variations have a red spectrum.

B. The loss of sensitivity due to a wrong PSD

Suppose that instead of the true PSD, $S_c(f)$, we use a wrong one due to PSD misestimation:

$$S_w(f) = S_c(f) (1 + \epsilon(f)). \quad (19)$$

In this section, we compute the bias to the recovered SNR of a putative event due to the wrong PSD. We will show that:

1. The first-order effect of using a wrong PSD in computing Eq. (9) is misestimation of its standard deviation.
2. We show that if the standard deviation of the overlaps is corrected, then the SNR loss will be of order $\epsilon^2(f)$.

Let us consider the following statistical model for the strain data:

$$d(f) = \alpha h(f) + n(f), \quad (20)$$

where $h(f)$ is a compact binary coalescence template waveform with an amplitude normalization α , and $n(f)$ is the noise which obeys the true PSD $S_c(f)$. For simple notation, we define

$$I(f) \equiv |h(f)|^2 / S_c(f), \quad (21)$$

which is normalized according to

$$\sum_f I(f) = 1. \quad (22)$$

We now compute the overlap using the wrong PSD:

$$z_w = \sum_f \frac{h^*(f) d(f)}{S_w(f)}. \quad (23)$$

In the absence of a signal $\alpha = 0$, this overlap has a variance λ_w :

$$\begin{aligned} \lambda_w &\equiv \langle |z_w|^2 \rangle \\ &= \sum_f \frac{h^*(f) h(f)}{S_w^2(f)} S_c(f) \\ &= \sum_f \frac{I(f)}{(1 + \epsilon(f))^2}. \end{aligned} \quad (24)$$

We note that this is the true variance of z_w , which is not easily computable as it depends on the unknown S_c . We will however measure it directly from data.

The naive way to calculate the variance involves using the wrong PSD S_w :

$$\begin{aligned} \lambda_{w, \text{computed}} &\equiv \sum_f \frac{h^*(f) h(f)}{S_w(f)} \\ &= \sum_f \frac{I(f)}{1 + \epsilon(f)}. \end{aligned} \quad (25)$$

In the presence of a signal, the expected response is given by:

$$\alpha_w \equiv \langle z_w \rangle_h = \alpha \sum_f \frac{|h(f)|^2}{S_w(f)} = \alpha \sum_f \frac{I(f)}{1 + \epsilon(f)} \quad (26)$$

Thus, if we use z_w as our statistic, the significance squared is

$$\rho_w^2 = \frac{\langle z_w \rangle_h^2}{\langle |z_w|^2 \rangle} = \frac{\alpha_w^2}{\lambda_w}. \quad (27)$$

If we were to take for granted that the variance of z_w was $\lambda_{w, \text{computed}}$, we would have estimated:

$$\rho_{w, \text{computed}}^2 = \frac{\alpha_w^2}{\lambda_{w, \text{computed}}}. \quad (28)$$

If the true PSD is known, the optimal overlap is given by:

$$z_c = \sum_f \frac{h^*(f) d(f)}{S_c(f)} \quad (29)$$

which maximizes the SNR. In the presence of a signal:

$$\alpha_c \equiv \langle z_c \rangle_h = \alpha \sum_f \frac{|h(f)|^2}{S_c(f)} = \alpha \sum_f I(f) = \alpha. \quad (30)$$

Absent the signal, z_c has a variance

$$\lambda_c \equiv \langle |z_c|^2 \rangle = \sum_f \frac{|h(f)|^2}{S_c^2(f)} S_c(f) = \sum_f I(f) = 1. \quad (31)$$

The optimal ρ_c^2 is then:

$$\rho_c^2 = \frac{\langle z_c \rangle_h^2}{\langle |z_c|^2 \rangle} = \frac{\alpha_c^2}{\lambda_c} = \alpha^2. \quad (32)$$

Therefore, the relative information loss between using the true PSD $S_c(f)$ and the wrong PSD $S_w(f)$ is given by:

$$\frac{\rho_c^2}{\rho_w^2} = \frac{\alpha_c^2 \lambda_w}{\alpha_w^2 \lambda_c} = \frac{\sum_f I(f) / (1 + \epsilon(f))^2}{\left[\sum_f I(f) / (1 + \epsilon(f)) \right]^2}. \quad (33)$$

If $\epsilon(f)$ was frequency independent, then the ratio would be one, and ρ_w^2 would be as optimal as ρ_c^2 . More generally, $\rho_c^2 > \rho_w^2$ for any frequency dependent $\epsilon(f)$, so there is always a sensitivity loss.

We now Taylor expand (33) and keep terms up to second order in $\epsilon(f)$:

$$\begin{aligned} \frac{\rho_c^2}{\rho_w^2} &= \frac{1 - 2 \sum_f I(f) \epsilon(f) + 3 \sum_f I(f) \epsilon^2(f)}{\left[1 - \sum_f I(f) \epsilon(f) + \sum_f I(f) \epsilon^2(f) \right]^2} \\ &= 1 + \sum_f I(f) \epsilon^2(f) - \left(\sum_f I(f) \epsilon(f) \right)^2 \geq 1. \end{aligned} \quad (34)$$

Thus, the leading correction to ρ^2 is quadratic with $\epsilon(f)$. This is understandable from the fact that the overlap z_c is constructed to maximize the SNR, and thus loss of sensitivity due to an error in PSD estimation should be a quadratic function.

The relative error in PSD estimation in each frequency band of $S_c(f)$ is $0.7 N_{\text{seg}}^{-1/2}$, where N_{seg} is the number of segments used for measuring the PSD using the Welch method (the coefficient 0.7 is determined empirically). Therefore, as a first corollary, the fractional sensitivity loss due to the statistical error in PSD estimation is $0.5 N_{\text{seg}}^{-1}$.

A key point is that computing ρ_w^2 requires λ_w , which cannot be computed without the knowledge about S_c . Nevertheless, λ_w can be directly measured from the data. We now compute the fractional error in the standard deviation estimation:

$$\frac{\lambda_w}{\lambda_{w, \text{computed}}} = \frac{\sum_f I(f)/(1 + \epsilon(f))^2}{\sum_f I(f)/(1 + \epsilon(f))} \quad (35)$$

$$= 1 - \sum_f I(f) \epsilon(f) + \mathcal{O}[\epsilon^2(f)].$$

At the leading order, the error scales linearly with $\epsilon(f)$. Thus if one uses $\lambda_{w, \text{computed}}$ to compute the SNR of an event one makes a linear mistake. We can think of $\sum_f [I(f)/(1 + \epsilon(f))]$ as a random variable which adds $\mathcal{O}[\epsilon(f)]$ noise to the estimate of the trigger SNRs. As we will describe in the next sections, this creates a tail in the distribution of triggers which results in substantial sensitivity loss. In the next sections we will discuss how to measure λ_w . As a remedy to PSD misestimation, this mitigates the sensitivity loss from $\mathcal{O}[\epsilon(f)] \sim 10\%$ (See Fig 1) to order $\mathcal{O}[\epsilon^2(f)] \sim 1\%$.

C. PSD drift correction

We propose a practical method to mitigate the sensitivity loss due to fast variations in the PSD (to second order in the amplitude of the drift) by tracking the time-dependent variance of the calculated overlaps at few percent accuracy. Since this is a single number to be estimated from the data (contrary to a full PSD which consists of thousands of numbers), we can estimate it using very short segments of data. For a 2% relative error in the variance, we will need roughly 10s of data. In Appendix B we estimate the precision with which this variance can be determined depending on the length of the data segments used. The variance around any given time t is empirically computed in the time domain via

$$\lambda_w(t) \equiv \langle |z_w(t)|^2 \rangle \approx \frac{1}{N_a} \sum_{t'=t-N_a\Delta t/2}^{t+N_a\Delta t/2} |z_w(t')|^2 \quad (36)$$

$$\propto \sum_{t'=t-N_a\Delta t/2}^{t+N_a\Delta t/2} |h_w(t') \otimes d_w(t')|^2$$

$$= \sum_f \left| \frac{h^*(f) d(f)}{S_w(f)} \right|^2$$

$$= \sum_f \frac{|h^*(f)|^2 |d(f)|^2}{S_w^2(f)},$$

where N_a is the number of scores used in the average, and h_w and d_w are the template and data whitened with the wrong PSD S_w , respectively. The equality in the third line is due to Parseval's theorem. The normalization factor in the above formula depends on the choice of N_a and can be precisely worked out.

The approximate equality in (36) is due to the finite number of samples used, which results in an uncertainty proportional to $N_a^{-1/2}$. Exact equality does not hold also because the PSD may vary within the averaging timescale $N_a\Delta t$. We invite the readers to read Appendix B for a full derivation.

It is noteworthy that $\lambda_w(t)$ is independent of the phase of $h(f)$ as long as the waveform has its support within $N_a\Delta t$. More generally, a correction that is computed for a particular waveform $h_0(f) = A_0(f) e^{i\psi_0(f)}$ is valid for any other waveform $h(f) = A(f) e^{i\psi(f)}$ who shares the same amplitude profile $A(f) = A_0(f)$ and whose time domain support, after dechirping,

$$h_{\text{dechirped}}(f) = h(f) e^{-i\psi_0(f)} \quad (37)$$

is shorter than $N_a\Delta t$.

In choosing the timescale on which the moving average is computed for $\lambda_w(t_0)$, there are two mutually competing factors. One is that we prefer the timescale to be as short as possible in order to capture the temporal variation in the PSD as much as possible. Another is that a smaller timescale for averaging leads to a larger sampling uncertainty in the measurement of $\lambda_w(t)$. We adopt a compromise in which the timescale is chosen to have 2% sampling uncertainty on $\lambda_w(t)$ (typically 10s in advanced LIGO O2 data). We estimate that for this value, the error due to not capturing faster variability is comparable, albeit varies slightly with time and from one interferometer to another. Another important issue is to ensure that the correction is not influenced by loud and abrupt noise transients or real GW events. We achieve this by computing Eq. (36) using a safe mean, namely we only average over overlaps that do not exceed 4.38σ .

In Figs. 3 and 4 we show the dramatic impact on the trigger distribution from applying the PSD drift correction. In the figures we compare the histogram of the PSD-drift-corrected $|z_w|^2/\lambda_w$ statistic and that of the usual $|z_w|^2/\lambda_{w, \text{computed}}$. The relative importance of this correction and signal consistency vetoes depends on the waveform duration. In Fig. 3 we show the trigger distribution for our binary black hole bank BBH (0,0) [18], which contains relatively long waveforms in the chirp-mass range $3\text{--}5 M_\odot$, here the PSD drift correction is the most important correction. In Fig. 4 we use a bank with shorter waveforms, and it can be seen that the PSD drift correction is secondary in importance to signal consistency vetoes. Still, once those have been applied, the PSD drift correction substantially improves the sensitivity.




FIG. 3: Cumulative trigger rates in bank BBH (0,0) (binary black holes whose chirp mass falls in the range $3\text{--}5 M_\odot$ [18]) using the entirety of the O2 bulk data. The initial trigger distribution (blue) is the distribution of the maximum overlap obtained every second, computed after removal of bad data, as described in Ref. [13]. In green is shown the final cumulative trigger distribution, after correcting for PSD drift and applying signal consistency vetoes. In black we show the cumulative trigger distribution when PSD drift is unaccounted for. For reference we show in orange the trigger distribution when the signal consistency vetoes are not applied. While the PSD drift correction multiplying real events is on average one (see Fig. 1), the background distribution undergoes a dramatic change. This is because if left unaccounted for, the fluctuations in places with variance misestimation dominate the tail of the trigger distribution. This effect becomes more severe at higher trigger significance.

We quantify the volume improvement by comparing the incoherent double detector trigger distributions from the entirety of O2 in banks BBH 0 and BBH 3 before and after PSD drift correction. We further restrict to $\lambda_w/\lambda_{w,\text{computed}} < 1.4$, as large correction values often indicate big disturbances in the spectrogram and the associated seconds are often rejectable by other tests. We also restricted the results to places where both detectors had comparable response ($20 < (z^2/\lambda)_{\text{H,L}} < 50$), where

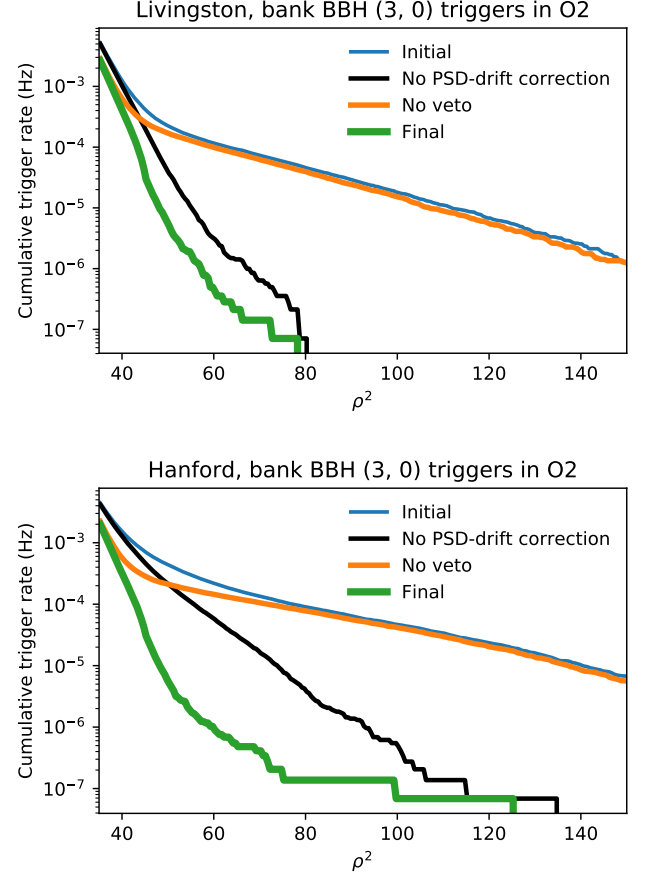


FIG. 4: Cumulative trigger rates in bank BBH (3,0) (binary black holes whose chirp mass falls in the range $20\text{--}40 M_\odot$ and whose effective spin does not have a very negative value [18]), using the entirety of the O2 bulk data. This bank contains most of the detected BBH events in O1 and O2. As can be seen, signal consistency checks for triggers are much more important in this domain as the waveforms in general have a very short duration in band. But by pushing back the background distribution, PSD drift correction still mitigates a substantial amount of sensitivity loss.

H, L refer to the LIGO detectors at Hanford and Livingston. We determine a 15% volume increase in both banks due to correcting the PSD drift effect. This is estimated using

$$\frac{V_{\text{corr}}}{V_{\text{uncorr}}} = \left[\frac{\left(\frac{z_H^2}{\lambda_H} + \frac{z_L^2}{\lambda_L} \right) \Big|_{\text{FAR}=1/\text{O2}}}{\left(\frac{z_H^2}{\lambda_{\text{H computed}}} + \frac{z_L^2}{\lambda_{\text{L computed}}} \right) \Big|_{\text{FAR}=1/\text{O2}}} \right]^{-3/2} \approx 1.15, \quad (38)$$

where the terms are evaluated at the threshold where the background distribution produces a false alarm rate of 1 trigger per O2 run. We expect this volume increase to be the same for BNSs as well as for NSBHs. This is a

conservative lower bound as many real events would not have comparable detector response, especially in the latest observing runs where the sensitivity differs between detectors by 20–40%. In the case where the detection practically hinges on one detector alone, the volume increase is 50% for Livingston and 150% for Hanford, as can be seen in Figs. 3 and 4. Since for a substantial amount of time in the second and third observing runs there is substantial asymmetry between the sensitivity of the detectors, as well as due to natural geometric considerations (and noise fluctuations), the volume contribution from this regime can be substantial.

III. HOLE FILLING

A. Signal processing rationale

The data acquired by gravitational wave detectors exhibits another non-Gaussian phenomenon that may hinder gravitational wave detection—“glitches” [22–24]. Glitches are abrupt disturbances of excess power, localized in time, and sometimes also in frequency. Such disturbances have diverse (and often unknown) physical or instrumental origins, and mathematical models that can be used to accurately characterize them are lacking. The timescale of disturbance ranges from a few milliseconds to a few seconds. The rate of glitches in the O2 data is time-variable, with an occurrence rate of about one per few hundred seconds. During searches for signals from compact binary coalescence, these disturbances induce candidate triggers that populate the tail of the distribution. Their presence dilutes the significance of genuine astrophysical triggers and degrades the search sensitivity. Removing these bad segments of data is not a trivial task, as simply zeroing them out (usually done using a “gate” that smoothly zeroizes the data) might result in **leakage of excess power which depending on other analysis choices may affect data for tens of seconds around and so be more harmful to the search effort than the disturbance itself**.

This is because spectral lines in the PSD whose inverse widths are longer than the duration of the gate leak power to neighboring seconds and frequency bins. See for example the middle panels of Figs. 5 and 6.

While very often these disturbance induced triggers are easily dismissed, sometimes careful treatment is necessary, as the time-domain footprint of real GW events (especially for BNS events whose waveforms last very long in band) may fortuitously overlap a disturbance. In that case, accurately determining the significance and estimating the merger parameters can be a complicated task. For example, GW170817 was in coincidence with a large disturbance in the Livingston detector (see Fig. 6). The analysis by the LIGO and Virgo Collaboration coherently removed this glitch in their analysis of GW170817 [6, 25], but most glitches lack an exact mathematical model. We therefore foresee that the analysis of future BNS detec-




FIG. 5: Effect of masking and inpainting glitches. *Top panel:* A segment^a of whitened strain data (in units of the noise standard deviation) that contains a glitch. The orange curve tracks the standard deviation σ calculated from a running window of 100 samples, and is typically close to unity as expected for whitened data. *Second panel:* Gating the glitch with an upside-down Tukey window (green) and then whitening generates artifacts in the whitened data, even outside the Tukey window. For example, σ stays above 1.1 for approximately 2 s to each side of the glitch (not shown). *Third panel:* The inpainted whitened data has unit variance outside the hole (shaded). *Bottom panel:* After inpainting, the “blued” strain is identically zero inside the hole, so overlaps with templates do not depend on the waveform information from inside the hole. Figure previewed already in the pipeline description paper, Ref. [13].

^a Data shown is from the Livingston observatory, GPS time: 1134155453.8789, first advanced LIGO observing run.

tions, especially as the sensitivity at low frequencies improves, will necessitate a treatment independent of any exact glitch model. In this section, we derive a method to exactly remove bad data segments, ensuring that the significance and inferred parameters of the event are not influenced by the offending segment of the data. We note that our methods to detect glitches, and their separation from genuine gravitational waves are described in full in [13]. In this paper we simply assume that the times of “bad data segments” that should be ignored are known.

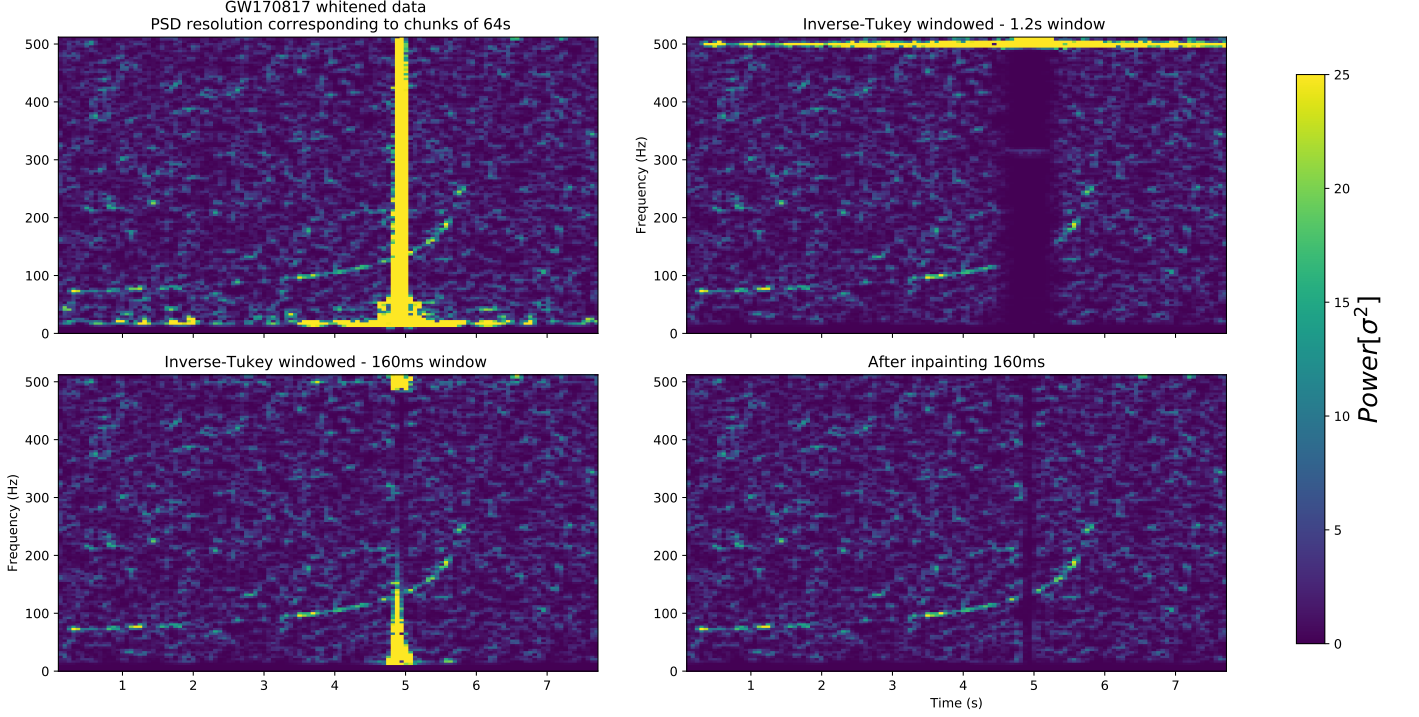


FIG. 6: Demonstration of inpainting in the segment of data containing GW170817. *Upper left panel:* whitened spectrogram^a of the original Livingston strain few seconds around the merger of GW170817, with a PSD measured using the welch method, with 64s chunks. Gating whitened data as in, e.g., Refs. [10, 11] would result in the shown power leakage in low frequencies ~ 4 seconds either side, caused by the interaction of the glitch with the whitening filter. *Upper right panel:* glitch gated away with an inverse Tukey window, with a timescale of 1.2s and $\alpha = 5/6$, as done in Ref. [6] Note the leakage from the spectral lines (around 500 Hz) is strongly affecting the data 5s to either side. *Bottom left panel:* Same as upper right, but with a 160ms window, leaving a more prominent noise leakage. *Bottom panel:* 160 ms inpainted by the hole-filling method presented in this work. The resulting data is free from disturbances and allows accurate inference and significance assessment. Note, the PSD resolution chosen in the making of this plot (reciprocal of the window size of 64s) differs from the one in Refs. [10, 12] (4s or 8s) (but is the default choice in Ref. [9]). We show similar plots with lower PSD resolution in Appendix C.

^a Spectrogram was made with 4Hz frequency resolution, 75% overlap between adjacent spectra.

B. Derivation

When computing overlaps one often assumes that the noise covariance is diagonal in the frequency domain and writes:

$$z_w = \sum_f \frac{h^*(f) d(f)}{S_w(f)}. \quad (39)$$

In practice, the data contain bad seconds that we need to mask out. Let us consider a data series of N_d samples and denote it as a vector with components d_i . We will denote the $N_d \times N_d$ covariance matrix of the noise as C_{ij} , which is diagonal in the frequency domain. Adopting the notation of linear algebra, the overlap can be cast into the matrix form:

$$z_w = C_{ij}^{-1} h_i^* d_j = h^\dagger C^{-1} d, \quad (40)$$

where h_i are the components of the template waveform and similarly for d_i .

In the presence of loud disturbances, we want the computed overlap to be independent of the behavior of the template waveform within the bad seconds. Let us assume that there is a list of samples in the time domain of length N_h to be masked. We denote $u^{(\alpha)}$ a list of N_h vectors such that each of these vectors is zero everywhere except at one of the samples to be masked. We can define the matrix A of size $N_d \times N_h$:

$$A_{i,\alpha} = u_i^{(\alpha)}, \quad (41)$$

and the matrix M of size $N_h \times N_h$ as:

$$\begin{aligned} M_{\alpha,\beta} &= C_{ij}^{-1} A_{i,\alpha} A_{j,\beta} \\ &= C_{ij}^{-1} u_i^{(\alpha)} u_j^{(\beta)} \\ &= A^T C^{-1} A. \end{aligned} \quad (42)$$

We are using the convention that Greek indices run over $1, \dots, N_h$ and roman indices run over $1, \dots, N_d$. We will now define the inpainting filter:

$$F = 1 - A M^{-1} A^T C^{-1} \quad (43)$$

and compute the scores with the inpainted data:

$$\tilde{z}_w = h^\dagger C^{-1} F d. \quad (44)$$

The presence of the hole changes the normalization of the template, so we renormalize \tilde{z} and compute:

$$z_w = \left(\frac{h^\dagger C^{-1} h}{h^\dagger C^{-1} F h} \right)^{1/2} \tilde{z}_w. \quad (45)$$

Note that this normalization is time and template dependent, which makes it computationally intractable to compute exactly everywhere. In the next subsection, we show how to approximate it using fast Fourier transform (FFT), via the stationary phase approximation.

The inpainting filter F has several desirable properties. First the score does not depend on the values of the template inside the hole because:

$$u^{(\alpha)T} C^{-1} F d = 0, \quad (46)$$

for any d . The inpainting filter also satisfies:

$$F^2 = F \quad (47)$$

$$F^T C^{-1} F = C^{-1} F = F^T C^{-1}. \quad (48)$$

Note that the filter F accomplishes its task by filling the hole with appropriate values. Values outside the holes are left untouched. The term $A M^{-1} A^T C^{-1}$ first weights the data by the inverse of the covariance. Then takes the values inside the holes (by multiplying by A^T) and makes some linear combinations (by multiplying them by M^{-1}). These values are then put back into the hole by multiplying these linear combinations by A . Note that the values outside of the holes are left untouched.

This hole filling procedure ensures that the scores do not depend on what the template does inside the hole. The only requirement is that the covariance matrix used to compute the scores (C^{-1}) is the same as the one used to build the matrix M .

We can obtain the same inpainting filter by considering a related problem. Suppose that the strain data d is the sum of a Gaussian random field h with covariance matrix C plus some additional source of noise n with covariance matrix N , $d = h + n$. Then the probability of h given the data d , $P(h|d)$, is:

$$P(h|d) \propto e^{-\frac{1}{2} (d-h)^\dagger N^{-1} (d-h)} e^{-\frac{1}{2} h^\dagger C^{-1} h}. \quad (49)$$

In the limit where the additional noise is zero outside the holes and infinite inside, the maximum of this function is given by $h = F d$.

Another equivalent formulation is that one is trying to find the maximum of

$$\chi^2 = \frac{1}{2} h^\dagger C^{-1} h, \quad (50)$$

with $h = d + \sum_\alpha a_\alpha u^{(\alpha)}$ for any a_α . That is, find an h that equals the data outside the holes but can take any value inside. This is achieved by taking $h = F d$.

C. Practical issues with inpainting

Correcting the variance

In order to assess the significance of overlaps computed on hole-filled data, we need to renormalize them using

$$\lambda_{\text{hole}} \equiv \frac{h^\dagger C^{-1} F h}{h^\dagger C^{-1} h}. \quad (51)$$

Under the stationary phase approximation, this coefficient could be understood as the fraction of SNR^2 that remains after the segment with bad data has been removed, i.e.,

$$\lambda_{\text{hole}}(t_0, h) \approx \frac{\sum_{t \notin \text{hole}} |h_w(t - t_0)|^2}{\sum_t |h_w(t)|^2}, \quad (52)$$

where t_0 denotes the merger time and h_w is the whitened waveform. Computing this for all times is then done via:

$$\lambda_{\text{hole}}(t_0, h) \approx \frac{(|\hat{h}_w|^2 \otimes \mathbb{1}_{\text{valid}})(t_0)}{\sum_t |h_w(t)|^2}, \quad (53)$$

where $\mathbb{1}_{\text{valid}}$ is one whenever the data is valid and zero otherwise. The convolution operation is computed via FFT using the convolution theorem.

In principle, the hole correction depends on the merger phase, but under the stationary phase approximation, and the approximation that the waveform would complete many cycles inside the hole, this correction would be the same for waveforms with different phases, and could be computed the same way for complex waveforms. The stationary phase approximation does not apply to short waveforms (say waveforms that are shorter than a few seconds), for these we do not compute the correction and we simply ignore overlaps that are too close to a hole. For short waveforms the correction is also not relevant as the fraction of data that is removed this way is smaller than a few percent of the total coincident time.

Edge holes

When data stream begins after a long break, or a long period of bad data, it is desirable to be able to detect events that are only partially inside the valid data. In order to do this optimally, we treat the edges of the data as holes, and fill them using the same formulae.

Reducing the computational complexity of hole filling

Since computing $F d$ requires inverting a matrix of size $N_h \times N_h$, which costs N_h^3 operations, it is of importance to find efficient ways of computing those. Since the whitening filter is of finite size (fixed in our pipeline to 64s), filling holes that are 64s apart could be done independently.

Hole filling a contiguous segment could be done faster via Toeplitz linear equation solution. This can be done through solving Eq. (46) for the vector Fd , using the fact that C^{-1} is Toeplitz, and that the only disturbed values are inside the hole. Segments of data with inseparable holes that have more than 10s of hole duration are assumed to be all bad and filled as contiguous segments using the Toeplitz solver. When the size of the hole is bigger than the size of the blueing filter (i.e., the whitening filter applied twice) we simply treat this as an edge, and fill only the duration of the blueing filter.

IV. SUMMARY

We have presented methods to deal with two kinds of non-Gaussianity in the strain data, one kind is the fast changes of the PSD, and the other is times of bad data containing abrupt noise transients.

We proposed computing a running estimate for the variance of the overlaps. We call this correction the PSD drift correction and we have shown that it correctly accounts for the first order errors due to having a wrong PSD in calculating the overlaps. We have shown that the error in estimating the overlap variance is of order 2% on 15s of data, a segment that is short enough so that error due to even faster PSD variations is of similar magnitude. Albeit the PSD drift effect in principle has complex frequency structure, the correction could be done via a simple robust running mean estimate for the local variance of the overlaps.

We have shown that applying this correction dramatically reduces the background in single-detector overlap distributions, and because of this the search volume is increased by at least 15%. This improvement is more substantial if the detection significance mainly hinges on one detector, as often happens due to the different antenna patterns of the detectors or their respective sensitivities. We point out that finding the physical source of

PSD change or better tracking of the fast changes could further improve the search sensitivity by 2 – 3%.

We have demonstrated a method to inpaint the strain data in unusable data segments. We have presented how this method allows to ignore loud strain artifacts, compute exact significance for genuine GW events that step into these loud artifacts and estimate the exact significance of such candidates. We expect all these methods to be instrumental in GW search in current and future data analysis.

ACKNOWLEDGMENTS

We greatly thank the LIGO Collaboration and the Virgo Collaboration for making the O1 and O2 data publicly accessible and easily usable.

This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center (<https://www.gw-openscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

BZ acknowledges the support of the Peter Svennilsson Membership Fund. TV acknowledges support by the Friends of the Institute for Advanced Study. TV also thanks the International Centre for Theoretical Sciences (ICTS) for hosting him during the period in which this manuscript was written. LD acknowledges the support from the Raymond and Beverly Sackler Foundation Fund. MZ is supported by NSF grants AST-1409709, PHY-1521097 and PHY-1820775 the Canadian Institute for Advanced Research (CIFAR) Program on Gravity and the Extreme Universe and the Simons Foundation Modern Inflationary Cosmology initiative.

Appendix A: Statistical properties of non-stationary noise

1. Covariance between Fourier components

In this section, we derive the correction to the definition of the discrete PSD, i.e., Eq. (5) in the presence of slow variations in the properties of the noise. The starting point are Eqs. (13) and (14), which define the behavior of the noise over short and long timescales.

As in Section II A, we work with DFT coefficients defined over a segment of data of length N . We are interested in the correlation between the DFT coefficients, which we expand as

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = \sum_{a,b} \langle d(a\Delta t) d(b\Delta t) \rangle e^{2\pi i(f_{m'}b - f_m a)\Delta t}. \quad (\text{A1})$$

Let us first work through the stationary case, where the expectation value on the right-hand side equals

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = \sum_{a,b} C_n [(b-a)\Delta t] e^{2\pi i(f_{m'}b - f_m a)\Delta t}. \quad (\text{A2})$$

Substituting the inverse of Eq. (3), we have

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = \sum_{a,b} \int_{-\infty}^{\infty} df S_{n,2}(f) e^{2\pi i f(b-a)\Delta t} e^{2\pi i(f_{m'}b - f_m a)\Delta t} \quad (\text{A3})$$

$$= \int_{-\infty}^{\infty} df S_{n,2}(f) \sum_{b=0}^N e^{2\pi i(f+f_{m'})b\Delta t} \sum_{a=0}^N e^{-2\pi i(f+f_m)a\Delta t}. \quad (\text{A4})$$

In the above equations, indices a and b run from 0 to $N-1$ in steps of 1 (as in the definition of the DFT in Footnote 1). Without loss of generality, we can shift the time-axis so that they run from $-N/2+1$ to $N/2$ instead. Evaluating the sums on the right-hand side of the above equation, we get

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = e^{\pi i(f_{m'} - f_m)\Delta t} \int_{-\infty}^{\infty} df S_{n,2}(f) W_N(f; f_m) W_N(f; f_{m'}), \quad (\text{A5})$$

where $W_N(f; f_m)$ is the following window function:

$$W_N(f; f_m) = \frac{\sin[\pi(f + f_m)N\Delta t]}{\sin[\pi(f + f_m)\Delta t]}. \quad (\text{A6})$$

As a function of frequency f , $W_N(f; f_m)$ has a series of peaks of amplitude N and alternating signs (for even N), and width $\sim 1/(N\Delta t)$, that pick out frequencies where the sines in the numerator and denominator vanish. The peak-frequencies are separated by the sampling frequency, so if we bandpassed the data before sampling, we can restrict to the primary frequency interval between $-f_s/2$ and $f_s/2$ (where $f_s = 1/\Delta t$ is the sampling frequency). If the PSD behaves smoothly on frequency scales of $\sim 1/(N\Delta t)$, $W_N(f; f_m)$ behaves like a delta-function selecting the appropriate frequency in the integrand, and we have

$$\left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle \approx \frac{N}{\Delta t} S_{n,2}(f_m) \delta_{m,m'} = \frac{N}{2\Delta t} S_n(f_m) \delta_{m,m'}. \quad (\text{A7})$$

Applying the whitening filter, we obtain

$$\left\langle \widetilde{d_w}(f_m) \left[\widetilde{d_w}(f_{m'}) \right]^* \right\rangle = N \delta_{m,m'}. \quad (\text{A8})$$

Next, we now include the non-stationary terms (the extra terms in Eqs. (13) and (14)) when simplifying Eq. (A1). They lead to an extra term, which we evaluate as follows:

$$\delta \left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = \sum_{a,b} \delta C_n \left[(b-a)\Delta t; \frac{a+b}{2}\Delta t \right] e^{2\pi i(f_{m'}b - f_m a)\Delta t} \quad (\text{A9})$$

$$= \sum_{a,b} \int_{-\infty}^{\infty} df S_{n,2}(f) \epsilon \left[f; \frac{a+b}{2}\Delta t \right] e^{2\pi i f(b-a)\Delta t} e^{2\pi i(f_{m'}b - f_m a)\Delta t}, \quad (\text{A10})$$

where in the second equation, we used the inverse of Eq. (15). It is convenient to define the variables A and B :

$$A = \frac{a+b}{2} \in \left\{ -\frac{N}{2} + 1, -\frac{N}{2} + \frac{3}{2}, \dots, \frac{N}{2} \right\}, \text{ and} \quad (\text{A11})$$

$$B = b-a \in \{-L_A, L_A+2, \dots, L_A\}, \text{ where} \quad (\text{A12})$$

$$L_A = \begin{cases} N-2+2A, & A \leq \frac{1}{2}, \\ N-2A, & A > \frac{1}{2}. \end{cases} \quad (\text{A13})$$

We recast Eq. (A10) in terms of these variables and simplify as follows:

$$\delta \left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle = \sum_A \sum_{B=-L_A}^{L_A} \int_{-\infty}^{\infty} df S_{n,2}(f) \epsilon(f; A\Delta t) e^{2\pi i f B \Delta t} e^{2\pi i[f_{m'}(A+B/2) - f_m(A-B/2)]\Delta t} \quad (\text{A14})$$

$$= \sum_A \int_{-\infty}^{\infty} df S_{n,2}(f) \epsilon(f; A\Delta t) e^{2\pi i(f_{m'} - f_m)A\Delta t} \sum_{B=-L_A}^{L_A} e^{2\pi i[f + (f_{m'} + f_m)/2]B\Delta t} \quad (\text{A15})$$

$$= \sum_A \int_{-\infty}^{\infty} df S_{n,2}(f) \epsilon(f; A\Delta t) e^{2\pi i(f_{m'} - f_m)A\Delta t} W_{L_A+1}(2f; f_m + f_{m'}) . \quad (\text{A16})$$

As earlier, we restrict to the primary frequency interval between $-f_s/2$ and $f_s/2$, within which the window function picks out two peaks at $f = -\bar{f}$ and $f = \pm f_s/2 - \bar{f}$ (here, $\bar{f} = (f_m + f_{m'})/2$, and in the second equation, we pick the sign that brings the right-hand side into the primary interval). The peaks have the same (opposite) signs when L_A is even (odd). The peaks are smeared out when L_A is of order unity, which happens only near the edges of the interval. We neglect these edge effects and replace the window function with the appropriate delta functions at its peaks:

$$\delta \left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle \approx \sum_A \int_{-\infty}^{\infty} df S_{n,2}(f) \epsilon(f; A\Delta t) e^{2\pi i(f_{m'} - f_m)A\Delta t} \frac{1}{2\Delta t} \times \begin{cases} \left[\delta(f + \bar{f}) - \delta\left(f + \bar{f} \pm \frac{f_s}{2}\right) \right] & A \in \left\{ -\frac{N}{2} + \frac{3}{2}, -\frac{N}{2} + \frac{5}{2}, \dots, \frac{N}{2} - \frac{1}{2} \right\}, \\ \left[\delta(f + \bar{f}) + \delta\left(f + \bar{f} \pm \frac{f_s}{2}\right) \right] & A \in \left\{ -\frac{N}{2} + 1, -\frac{N}{2} + 2, \dots, \frac{N}{2} \right\} \end{cases} . \quad (\text{A17})$$

Let us denote the frequency offset $f_m - f_{m'}$ by Δf . The contributions of the second term inside the square brackets to the sum approximately cancel (since the parts with alternating signs are approximately equal), and hence we can simplify the sum to

$$\delta \left\langle \tilde{d}(f_m) \left[\tilde{d}(f_{m'}) \right]^* \right\rangle \approx \frac{1}{2\Delta t} \sum_A S_{n,2}(-\bar{f}) \epsilon(-\bar{f}; A\Delta t) e^{-2\pi i \Delta f A\Delta t} \quad (\text{A18})$$

$$\approx \frac{S_n(\bar{f})}{2\Delta t} \tilde{\epsilon}(\bar{f}; \Delta f), \quad (\text{A19})$$

where $\tilde{\epsilon}(\bar{f}; \Delta f)$ is the DFT, evaluated at the frequency offset Δf , of $\epsilon(\bar{f}, T = n\Delta t)$ sampled at a rate of $f_s = 1/(\Delta t)$.³ Finally, we combine Eqs. (A7) and (A19) and apply the whitening filter to obtain the generalized version of Eq. (7):

$$\left\langle \widetilde{d_w}(f_m) \left[\widetilde{d_w}(f_{m'}) \right]^* \right\rangle \approx N\delta_{m,m'} + \frac{S_n(\bar{f})}{[S_n(f_m) S_n(f_{m'})]^{1/2}} \tilde{\epsilon}(\bar{f}; \Delta f). \quad (\text{A20})$$

We assume that (a) the changes to the PSD $S_n(f)$ occur on long timescales, i.e., $\Delta f \ll \bar{f}$, and (b) $S_n(f)$ is smooth on frequency scales $\sim \Delta f$. Then, we can write

$$\left\langle \widetilde{d_w}(f_m) \left[\widetilde{d_w}(f_{m'}) \right]^* \right\rangle \approx N\delta_{m,m'} + \tilde{\epsilon}(\bar{f}; \Delta f). \quad (\text{A21})$$

In our application, the first assumption above holds very well, while the second assumption fails only in the immediate vicinity of spectral lines.

2. Power spectrum of the variance of the matched filtering scores

In this section, we consider the variance of the matched filtering scores $z(t)$ (as defined in Eq. (11)) as a time series, and derive its PSD. We empirically define the variance as the output of the estimator:

$$\langle z^2 \rangle(t = n\Delta t) = \sum_{n'} z^2[t = (n + n')\Delta t] w(-n'\Delta t), \quad (\text{A22})$$

where $\sum_n w(n\Delta t) = 1$. The window function w effectively manifests as a multiplicative low-pass filter in the formulae in the rest of this section, and does not change any details below its cutoff frequency. For ease of presentation, we will omit the window function and directly use the power, z^2 , in place of the variance $\langle z^2 \rangle$.

The DFT of the power is

$$\widetilde{z^2}(f_m) = \frac{1}{N} \sum_p \left[\widetilde{h_w}(f_p) \right]^* \widetilde{d_w}(f_p) \widetilde{h_w}(f_p - f_m) \left[\widetilde{d_w}(f_p - f_m) \right]^*. \quad (\text{A23})$$

³ Note that the sum in Eq. (A18) is on a finer grid sampled at frequency $2/(\Delta t)$. Equation (A19) holds assuming that ϵ is band-

limited below $f_s = 1/(\Delta t)$, which is true in our case.

In the above equation, the convention is that frequencies outside the primary interval of $(-f_s/2, f_s/2]$ are replaced by the values in the interval that they alias to. We view the power as a random time series, with its own autocorrelation function. The PSD of this random series is defined by:

$$\langle \tilde{z}^2(f_m) [\tilde{z}^2(f_{m'})]^* \rangle = \frac{N}{2\Delta t} S_{z^2}(f_m) \delta_{m,m'}. \quad (\text{A24})$$

We substitute Eq. (A23) and obtain:

$$\begin{aligned} \langle \tilde{z}^2(f_m) [\tilde{z}^2(f_{m'})]^* \rangle &= \frac{1}{N^2} \sum_{p,q} [\widetilde{h_w}(f_p)]^* \widetilde{h_w}(f_p - f_m) \widetilde{h_w}(f_q) [\widetilde{h_w}(f_q - f_{m'})]^* \times \\ &\quad \langle \widetilde{d_w}(f_p) [\widetilde{d_w}(f_p - f_m)]^* [\widetilde{d_w}(f_q)]^* \widetilde{d_w}(f_q - f_{m'}) \rangle. \end{aligned} \quad (\text{A25})$$

We first work through the case with stationary noise. We use Wick's theorem to simplify the expectation value on the right-hand side as the sum of pairwise products, which we evaluate using Eq. (A8):

$$\langle \tilde{z}^2(f_m) [\tilde{z}^2(f_{m'})]^* \rangle = \sum_{p,q} [\widetilde{h_w}(f_p)]^* \widetilde{h_w}(f_p - f_m) \widetilde{h_w}(f_q) [\widetilde{h_w}(f_q - f_{m'})]^* [\delta_{m,0}\delta_{m',0} + \delta_{p,q}\delta_{m,m'} + \delta_{p+q,m}\delta_{m,m'}]. \quad (\text{A26})$$

We obtain an expression for the PSD of the power, S_{z^2} , by substituting Eq. (A24) and simplifying:

$$S_{z^2}(f_m) = \frac{2\Delta t}{N} \sum_{p,q} [\widetilde{h_w}(f_p)]^* \widetilde{h_w}(f_p - f_m) \widetilde{h_w}(f_q) [\widetilde{h_w}(f_q - f_m)]^* [\delta_{m,0} + \delta_{p,q} + \delta_{p+q,m}] \quad (\text{A27})$$

$$= 2N\Delta t \delta_{m,0} + \frac{4\Delta t}{N} \sum_{p,q} |\widetilde{h_w}(f_p)|^2 |\widetilde{h_w}(f_q)|^2 \delta_{p+q,m}. \quad (\text{A28})$$

$$= 2N\Delta t \delta_{m,0} + 4\Delta t \left| \widetilde{h_w \otimes \overleftarrow{h_w}} \right|^2(f_m). \quad (\text{A29})$$

In the final equation, $\widetilde{h_w \otimes \overleftarrow{h_w}}$ is the autocorrelation function of the whitened waveform. To simplify the first term on the right-hand side, we used the fact that the SNR in Eq. (10) is normalized to unity.

In the non-stationary case, the PSD of the power, S_{z^2} , also depends on the varying part of the noise PSD, $\epsilon(f, T)$. Given a number of frequency-bins, f_m (conjugate to the 'fast' timescales), we consider the $\epsilon(f_m, T)$ as a set of random time-series, varying over long timescales T . The most general form of the correlations between these time-series is

$$\langle \tilde{\epsilon}(f_m, f_a) [\tilde{\epsilon}(f_{m'}, f_b)]^* \rangle = \frac{N}{2\Delta t} S_{\epsilon, (m, m')}(f_a) \delta_{a,b}. \quad (\text{A30})$$

In the above equation, $\tilde{\epsilon}(f_m, f_a)$ represents the DFT of $\epsilon(f_m, T = n\Delta t)$ sampled at a rate $f_s = 1/(\Delta t)$. If we perform the Singular Value Decomposition (SVD) of the matrix $S_{\epsilon, (m, m')}(f_a)$, the basis vectors represent frequency components (with weights given by the coefficients) that vary together. We assume the simplest possible form of the correlations in Eq. (A30):

$$\langle \tilde{\epsilon}(f_m, f_a) [\tilde{\epsilon}(f_{m'}, f_b)]^* \rangle = \frac{N}{2\Delta t} S_{\epsilon}(f_a) \delta_{a,b}. \quad (\text{A31})$$

This implies that all the $\epsilon(f_m, T)$ vary in step with each other, and with the same amplitude.

We return to Eq. (A25) to evaluate the PSD of the power, S_{z^2} . We first fix a particular realization of $\epsilon(f_m, T)$, compute S_{z^2} using the disconnected terms in Eq. (A21), and additionally average over the realizations of ϵ using Eq. (A31) (we neglect the connected part of the expectation value in Eq. (A21)). In addition to the terms in Eq. (A29), the following terms show up:

$$\begin{aligned} \delta S_{z^2}(f_m) &= \frac{1}{N^2} \sum_{p,q} [\widetilde{h_w}(f_p)]^* \widetilde{h_w}(f_p - f_m) \widetilde{h_w}(f_q) [\widetilde{h_w}(f_q - f_m)]^* [S_{\epsilon}(f_m) + S_{\epsilon}(f_p - f_q) + S_{\epsilon}(f_p + f_q - f_m)] \\ &= |\widetilde{h_w}(f_m)|^2 S_{\epsilon}(f_m) \\ &\quad + \frac{1}{N^2} \sum_{p,q} [\widetilde{h_w}(f_p)]^* \widetilde{h_w}(f_p - f_m) \widetilde{h_w}(f_q) [\widetilde{h_w}(f_q - f_m)]^* [S_{\epsilon}(f_p - f_q) + S_{\epsilon}(f_p + f_q - f_m)]. \end{aligned} \quad (\text{A32})$$

We can see by a change of variable that both the terms in the summation on the right-hand side of Eq. (A32) are equal. Carrying this through,

$$\begin{aligned}
\delta S_{z^2}(f_m) &= |\widetilde{h_w}^2(f_m)|^2 S_\epsilon(f_m) + \frac{2}{N^2} \sum_{p,q} \left[\widetilde{h_w}(f_p) \right]^* \widetilde{h_w}(f_p - f_m) \widetilde{h_w}(f_q) \left[\widetilde{h_w}(f_q - f_m) \right]^* S_\epsilon(f_p - f_q) \\
&= |\widetilde{h_w}^2(f_m)|^2 S_\epsilon(f_m) \\
&\quad + \frac{2}{N^2} \sum_{q, \Delta q \equiv (p-q)} \widetilde{h_w}(f_q) \left[\widetilde{h_w}(f_q + f_{\Delta q}) \right]^* \left[\widetilde{h_w}(f_q - f_m) \right]^* \widetilde{h_w}(f_q + f_{\Delta q} - f_m) S_\epsilon(f_{\Delta q}) \\
&= |\widetilde{h_w}^2(f_m)|^2 S_\epsilon(f_m) + \frac{2}{N} \sum_{\Delta q} \left| \widetilde{h_w} \circledast \widetilde{h_w} e^{-2\pi i f_{\Delta q} t} \right|^2 (f_m) S_\epsilon(f_{\Delta q}).
\end{aligned} \tag{A33}$$

In the second term on the right-hand side, we multiplied the signal with an oscillatory function of time ($e^{-2\pi i f_{\Delta q} t}$) before reversing and convolving it with itself.

We can simplify the form of Eq. (A33) using the following approximations: (a) variations in the noise PSD ($\tilde{\epsilon}$) have support at much lower frequencies than the whitened signal $\widetilde{h_w}$ does, and (b) the time-domain whitened waveform is much shorter than the timescales over which $\epsilon(f; T)$ varies. The first assumption holds very well, since $\epsilon(f; T)$ varies on timescales of tens of seconds (i.e., $\tilde{\epsilon}$ has power at $f \lesssim 0.1$ Hz) while signals accumulate SNR only at $f \gtrsim 20$ Hz; the second assumption does not hold for signals that are longer than a few tens of seconds (such as signals from merging binary neutron stars).

Under the second approximation above, the factor of $e^{-2\pi i f_{\Delta q} t}$ in Eq. (A33) acts like a constant phase multiplying the signal, which makes no difference. We can pull it out, and interpret the summation as a Riemann sum to obtain

$$\delta S_{z^2}(f_m) \approx \left| \widetilde{h_w}^2(f_m) \right|^2 S_\epsilon(f_m) + 2\Delta t \left| \widetilde{h_w} \circledast \widetilde{h_w} \right|^2 (f_m) \int_{-\infty}^{\infty} df S_\epsilon(f). \tag{A34}$$

We combine Eqs. (A29) and (A34) to get our final simplified form of the PSD of the variations in the power of the matched filtering scores:

$$S_{z^2}(f_m) \approx 2N\Delta t \delta_{m,0} + 2\Delta t \left| \widetilde{h_w} \circledast \widetilde{h_w} \right|^2 (f_m) \left[2 + \int_{-\infty}^{\infty} df S_\epsilon(f) \right] + \left| \widetilde{h_w}^2(f_m) \right|^2 S_\epsilon(f_m). \tag{A35}$$

Appendix B: Computation of the measurement error in the PSD drift

In this section, we work with whitened waveforms and data, $h_w(n\Delta t)$ and $d_w(n\Delta t)$ respectively, that are functions of the discrete index n , and were whitened using some fiducial filter. The auto-correlation function of perfectly white data is

$$\langle d_w(t) d_w(t + n\Delta t) \rangle = \delta_{n,0}, \tag{B1}$$

where the lowercase δ stands for the Kronecker delta. We define the time-series of matched filtering overlaps as

$$z(n\Delta t) = \sum_m h_w(m\Delta t) d_w[(n-m)\Delta t], \tag{B2}$$

which is equivalent to Eq. (11). We now consider the statistics of these overlaps.

Let us first assume the whitening was successful, i.e., Eq. (B1) applies. In this case, the autocorrelation function of the score $z(t)$ equals that of the waveform, i.e.,

$$\langle z(t) z(t + n\Delta t) \rangle = \left[h_w \circledast \widetilde{h_w} \right] (n\Delta t) \tag{B3}$$

In particular, if we normalize the whitened waveform to satisfy $\sum_n |h_w(n\Delta t)|^2 = 1$, the autocorrelation function equals unity at zero-lag, i.e., the score has unit variance. For typical waveforms, the autocorrelation function has a very small time-domain width (for the waveform for GW150914, the time-domain half-width, defined as the lag at which the autocorrelation function drops to 0.5, equals ~ 1.5 ms; at our fiducial sampling rate of 1024 Hz, ~ 2.5 samples are correlated).

If we fail to whiten the data perfectly, the scores $z(t)$ are still normally distributed, but with a different variance. Formally, the new auto-correlation function is

$$V(n\Delta t) = \langle z(t) z(t + n\Delta t) \rangle = \sum_{a,b} h_w(a\Delta t) h_w(b\Delta t) \langle d_w^*(t - a\Delta t) d_w[t + (n - b)\Delta t] \rangle. \quad (\text{B4})$$

We can write the convolutions in Eq. (B4) in the Fourier domain.

$$V(n\Delta t) = \frac{1}{N_a^2} \sum_{p,q} [\widetilde{h_w}(f_p)]^* \widetilde{h_w}(f_q) \left\langle [\widetilde{d_w}(f_p)]^* \widetilde{d_w}(f_q) \right\rangle e^{2\pi i(f_q - f_p)t} e^{2\pi i f_q \Delta t}, \quad (\text{B5})$$

where N_a is the number of samples used to compute the PSD drift correction. Instead of the general approach we followed in Section II, we use a segment of data that is shorter than the typical timescales over which the PSD varies, but assume that we used the wrong local PSD. In this case, the incorrectly whitened the data, d_w , is locally described by a stationary and real-valued Gaussian random variable:

$$\left\langle [\widetilde{d_w}(f_p)]^* \widetilde{d_w}(f_q) \right\rangle = \frac{N_a}{2\Delta t} S_{d_w}(f_p) \delta_{p,q}, \quad (\text{B6})$$

where S_d is the two-sided power spectral density of the process. Substituting in Eq. (B5) gives us

$$V(n\Delta t) = \frac{1}{N_a \Delta t} \sum_p \frac{1}{2} |\widetilde{h_w}(f_p)|^2 S_{d_w}(f_p) e^{2\pi i f_p n \Delta t}. \quad (\text{B7})$$

In the case where d_w is white noise, $S_{d_w} = 2\Delta t$, and using Parseval's identity, the zero-lag auto-correlation function of Eq. (B7) equals the normalization factor $\sum_n |h_w(n)|^2 = 1$.

In practice, we use the estimator

$$\mathcal{V} = \frac{1}{N_a^2} \sum_p |\widetilde{h_w}(f_p)|^2 |\widetilde{d_w}(f_p)|^2 = \frac{1}{N_a} \sum_n z^2(n\Delta t) \quad (\text{B8})$$

The variance of this estimator is

$$\begin{aligned} \sigma_{\mathcal{V}}^2 &= \langle \mathcal{V}^2 \rangle - \langle \mathcal{V} \rangle^2 \\ &= \frac{1}{N_a^4} \sum_{p,q} |\widetilde{h_w}(f_p)|^2 |\widetilde{h_w}(f_q)|^2 \left[\left\langle [\widetilde{d_w}(f_p)]^* \widetilde{d_w}(f_p) [\widetilde{d_w}(f_q)]^* \widetilde{d_w}(f_q) \right\rangle - \left\langle [\widetilde{d_w}(f_p)]^* \widetilde{d_w}(f_p) \right\rangle \left\langle [\widetilde{d_w}(f_q)]^* \widetilde{d_w}(f_q) \right\rangle \right] \\ &= \frac{1}{N_a^4} \sum_{p,q} |\widetilde{h_w}(f_p)|^2 |\widetilde{h_w}(f_q)|^2 \left[\left\langle [\widetilde{d_w}(f_p)]^* [\widetilde{d_w}(f_q)]^* \right\rangle \left\langle \widetilde{d_w}(f_p) \widetilde{d_w}(f_q) \right\rangle + \left\langle [\widetilde{d_w}(p)]^* \widetilde{d_w}(q) \right\rangle \left\langle \widetilde{d_w}(p) [\widetilde{d_w}(q)]^* \right\rangle \right] \\ &= \frac{1}{N_a^4} \sum_{p,q} |\widetilde{h_w}(f_p)|^2 |\widetilde{h_w}(f_q)|^2 \left(\frac{N}{2\Delta t} \right)^2 S_{d_w}^2(f_p) [\delta_{p,-q} + \delta_{p,q}] \\ &= \frac{1}{2(N_a \Delta t)^2} \sum_p |\widetilde{h_w}(f_p)|^4 S_{d_w}^2(f_p). \end{aligned} \quad (\text{B9})$$

We used Wick's theorem in going from the second to the third line, and the real-valued nature of the template and the data in the subsequent lines. If the data d_w were white, as earlier, we have $S_{d_w} = 2\Delta t$. Substituting this, we get

$$\sigma_{\mathcal{V}}^2 \Big|_{\text{white } d_w} = \frac{2}{N_a^2} \sum_p |\widetilde{h_w}(f_p)|^4 = \frac{4}{N_a^2} \sum_{p>0} |\widetilde{h_w}(f_p)|^4. \quad (\text{B10})$$

In the above formula, note that the FFT of the whitened template is performed on the interval of length N_a . In practice, we want to solve for N_a given a particular tolerance requirement. Suppose we have the whitened template on a frequency grid conjugate to a fiducial window of length N_0 . From the connection between the DFT and the continuous Fourier transform, we can scale the sum in Eq. (B10) according to the size of the grid.

$$\sigma_{\mathcal{V}}^2 \Big|_{\text{white } d_w} \approx \frac{4}{N_a N_0} \sum_{p_{N_0}>0} |\widetilde{h_w}(f_{p_{N_0}})|^4. \quad (\text{B11})$$

The formula above suggests that to measure the $(\lambda_w)^2 = \langle \mathcal{V} \rangle = V$ to two percent relative error, we need $\mathcal{O}(15\,000)$ samples, i.e., ~ 15 s for a typical waveform.

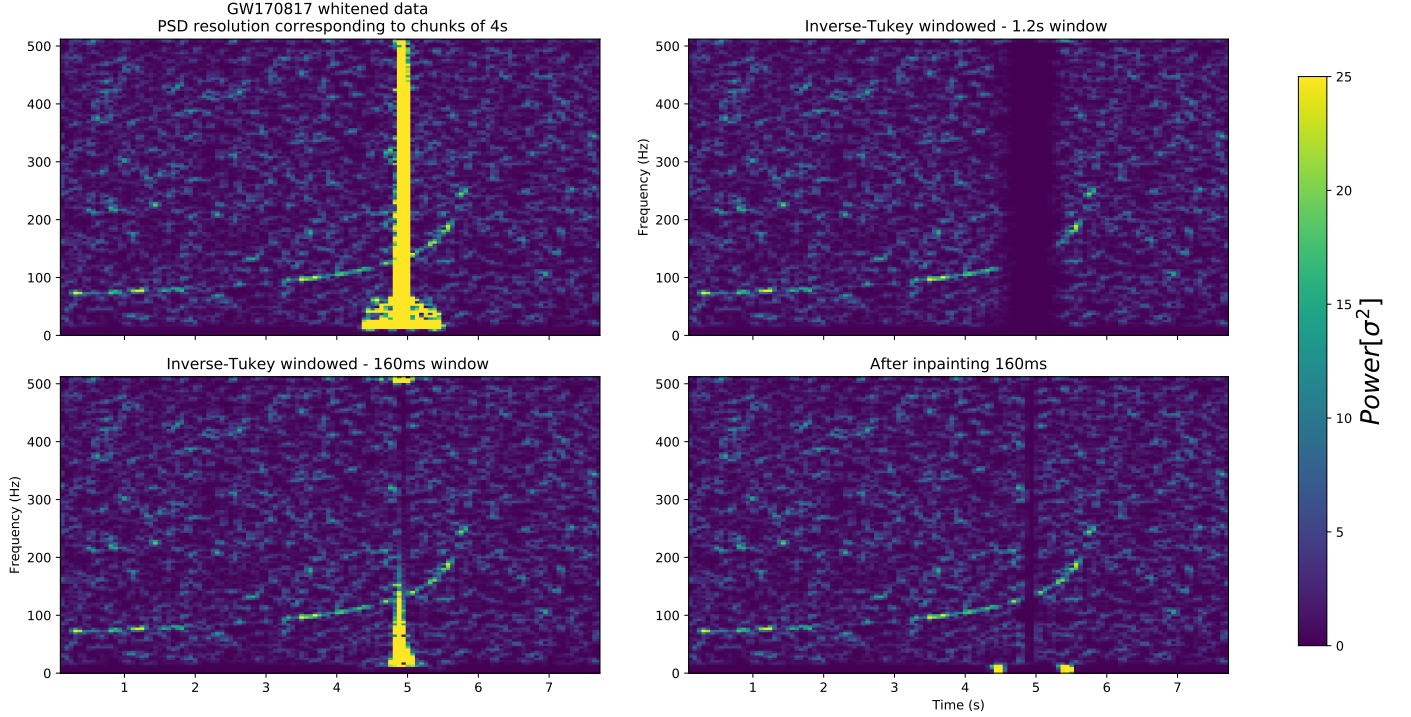


FIG. 7: Same as Figure 6 but with a PSD with a lower frequency resolution of 1/4 Hz. Prior to whitening, we had to manually notch a loud spectral line at 500 Hz, as it was unresolved and presented with a varying amplitude, which caused artifacts in the whitened data. It can be seen that at this frequency resolution, the gate that was used in the official analysis of GW170817 was enough to cover all the residuals from the glitch. It can also be seen that in order to make a narrower hole, inpainting is essential.

Appendix C: Gating whitened data with 4 s PSD resolutions

The whitened spectrogram in Figure 6 looks qualitatively different than the ones presented in the GW170817 discovery paper [6]. The reason is that we have used the standard PSD resolution of our pipeline (conjugate to 64s chunks) in its making. **Using a PSD resolution conjugate to 4s-16s, as done in [10, 12],** leads to a much more confined leakage ($\sim 1s$), see Figure 7. The reason is that using a lower PSD resolution, the spectral lines in the PSD are broader, leading to a short duration whitening filter⁴. In this configuration, the 1.2 s gate used in the official analysis of GW170817 does remove the entirety of the glitch. In any configuration (including this one), hole-filling is essential if we wish surgically excise only the 160ms containing the glitch. This ability is used extensively by our pipeline [13], in order to remove very short glitches, allowing the thorough cleaning of the data required for Gaussian dominated trigger distribution at very high S/N.

-
- [1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Physical Review Letters **116**, 061102 (2016).
 - [2] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Physical Review Letters **116**, 241103 (2016).

- [3] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. X **6**, 041015 (2016).
- [4] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **118**, 221101 (2017).
- [5] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), The Astrophysical Journal **851**, L35 (2017).
- [6] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Phys. Rev. Lett. **119**, 161101 (2017).

⁴ At this frequency resolution we had to notch the 500Hz line prior to filtering.

- (2017).
- [7] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. X* **9**, 031040 (2019).
 - [8] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, arXiv e-prints, arXiv:1902.10331 (2019), arXiv:1902.10331 [astro-ph.HE].
 - [9] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, (2019), arXiv:1904.07214v1 [astro-ph.HE].
 - [10] S. Sachdev, S. Caudill, H. Fong, R. K. L. Lo, C. Messick, D. Mukherjee, R. Magee, L. Tsukada, K. Blackburn, P. Brady, P. Brockill, K. Cannon, S. J. Chamberlin, D. Chatterjee, J. D. E. Creighton, P. Godwin, A. Gupta, C. Hanna, S. Kapadia, R. N. Lang, T. G. F. Li, D. Meacher, A. Pace, S. Privitera, L. Sadeghian, L. Wade, M. Wade, A. Weinstein, and S. L. Xiao, (2019), arXiv:1901.08580v1 [gr-qc].
 - [11] C. Messick, K. Blackburn, P. Brady, P. Brockill, K. Cannon, R. Cariou, S. Caudill, S. J. Chamberlin, J. D. E. Creighton, R. Everett, C. Hanna, D. Keppel, R. N. Lang, T. G. F. Li, D. Meacher, A. Nielsen, C. Pankow, S. Privitera, H. Qi, S. Sachdev, L. Sadeghian, L. Singer, E. G. Thomas, L. Wade, M. Wade, A. Weinstein, and K. Wiesner, *Phys. Rev. D* **95**, 042001 (2017), arXiv:1604.04324 [astro-ph.IM].
 - [12] S. A. Usman, A. H. Nitz, I. W. Harry, C. M. Biwer, D. A. Brown, M. Cabero, C. D. Capano, T. Dal Canton, T. Dent, S. Fairhurst, M. S. Kehl, D. Keppel, B. Krishnan, A. Lenon, A. Lundgren, A. B. Nielsen, L. P. Pekowsky, H. P. Pfeiffer, P. R. Saulson, M. West, and J. L. Willis, *Classical and Quantum Gravity* **33**, 215004 (2016), arXiv:1508.02357 [gr-qc].
 - [13] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, *Phys. Rev. D* **100**, 023011 (2019).
 - [14] B. P. Abbott *et al.*, *Classical and Quantum Gravity* **35**, 065010 (2018), arXiv:1710.02185 [gr-qc].
 - [15] “Gravitational Wave Open Science Center (GWOSC),” www.gw-openscience.org/02/ (accessed 2019-02-27).
 - [16] M. Vallisneri, J. Kanner, R. Williams, A. Weinstein, and B. Stephens, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 610 (2015) p. 012021, arXiv:1410.4839 [gr-qc].
 - [17] T. Venumadhav *et al.*, in preparation.
 - [18] J. Roulet, L. Dai, T. Venumadhav, B. Zackay, and M. Zaldarriaga, *Phys. Rev. D* **99**, 123022 (2019).
 - [19] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, *Phys. Rev. D* **85**, 122006 (2012), arXiv:gr-qc/0509116 [gr-qc].
 - [20] B. S. Sathyaprakash and S. V. Dhurandhar, *Phys. Rev. D* **44**, 3819 (1991).
 - [21] L. S. Finn, *Phys. Rev. D* **46**, 5236 (1992), arXiv:gr-qc/9209010 [gr-qc].
 - [22] L. Blackburn, L. Cadonati, S. Caride, S. Caudill, S. Chatterji, N. Christensen, J. Dalrymple, S. Desai, A. Di Credico, G. Ely, J. Garofoli, L. Goggin, G. González, R. Gouaty, C. Gray, A. Gretarsson, D. Hoak, T. Isogai, E. Katsavounidis, J. Kissel, S. Klimenko, R. A. Mercer, S. Mohapatra, S. Mukherjee, F. Raab, K. Riles, P. Saulson, R. Schofield, P. Shawhan, J. Slutsky, J. R. Smith, R. Stone, C. Vorvick, M. Zanolin, N. Zotov, and J. Zweizig, *Classical and Quantum Gravity* **25**, 184004 (2008), arXiv:0804.0800 [gr-qc].
 - [23] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Classical and Quantum Gravity* **33**, 134001 (2016), arXiv:1602.03844 [gr-qc].
 - [24] L. K. Nuttall, T. J. Massinger, J. Areeda, J. Betzwieser, S. Dwyer, A. Effler, R. P. Fisher, P. Fritschel, J. S. Kissel, A. P. Lundgren, D. M. Macleod, D. Martynov, J. McIver, A. Mullavey, D. Sigg, J. R. Smith, G. Vajente, A. R. Williamson, and C. C. Wipf, *Classical and Quantum Gravity* **32**, 245005 (2015), arXiv:1508.07316 [gr-qc].
 - [25] N. J. Cornish and T. B. Littenberg, *Classical and Quantum Gravity* **32**, 135012 (2015).