Mitigating the effect of glitches on gravitational-wave parameter estimation using an inpainting filter

Viviana A. Cáceres Barbosa
Physics Department, University of Puerto Rico at Mayagüez

Mentor: Derek Davis
LIGO, California Institute of Technology
(Dated: October 3, 2022)

Final Report
LIGO Caltech SURF Program 2022

Recovering accurate distributions for the source parameters of gravitational-wave signals is essential to confirm current models of general relativity and understand astrophysical properties of the universe. Glitches in gravitational-wave strain data may cause a bias in parameter estimation analyses that use Bayesian inference. We implement inpainting to address this problem in `Bilby`, one of various parameter estimation pipelines used for gravitational-wave analyses. Using two different methods to obtain inpainted data, we study how each process increases likelihood evaluation times. We also inject a binary black hole signal and study how inpainting with different configurations affects `Bilby`'s ability to recover accurate posterior distributions. The results suggest that our implementation exhibits the expected behavior of reducing the calculated likelihood without introducing a bias, and can now be added onto `Bilby` for more rigorous testing.

I. INTRODUCTION

Raw strain data recorded by gravitational-wave (GW) detectors such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) is typically dominated by noise coming from a variety of different sources. Although many of these sources are known and well-modeled, there are often short noise bursts, known as glitches, of unknown origin that impact the sensitivity of the data analysis softwares used by the LIGO-Virgo-KAGRA collective (LVK) [1]. Much of LVK efforts are dedicated to mitigating the effect of glitches in GW signal searches and parameter estimation (PE) analyses.

`Bilby` is one of various PE pipelines used by the LVK [2]. Like most PE pipelines, it uses Bayesian inference to produce posteriors, which are probability distributions of the GW source parameters. These are computed using Bayes' theorem:

$$p(\theta|d) = \frac{\mathcal{L}(d|\theta)\pi(\theta)}{z(d)} \tag{1}$$

where $\mathcal{L}(d|\theta)$ is the likelihood of measuring the data d given some source parameters θ, $\pi(\theta)$ is the prior distribution of these source parameters, and $z(d)$ is the evidence [3].

The `Bilby` analysis assumes that the noise in the data is stationary and Gaussian. This allows the use of the Gaussian noise likelihood [4]:

$$\mathcal{L}(d|\theta) = \frac{1}{|2\pi C|^{\frac{1}{2}}} \exp\left\{-\frac{1}{2}\chi^2(d,h)\right\} \tag{2}$$

with

$$\chi^2(d,h) = [d-h(\theta)]\, C^{-1}\, [d-h(\theta)] \tag{3}$$

where d is a vector representation of the data, C is the noise covariance matrix and $h(\theta)$ is the waveform that depends on parameters θ. In practice, Eq. 2 is costly to compute, so the Whittle approximation to the likelihood is used [5]:

$$\mathcal{L}(d|\theta) \propto \exp\left[-\frac{1}{2}(d-h|d-h)\right] \tag{4}$$

which expands to

$$\mathcal{L}(d|\theta) \propto \exp\left[-\frac{1}{2}(d|d) + (d|h) - \frac{1}{2}(h|h)\right] \tag{5}$$

where $(a|b)$ is the noise-weighted inner product defined as

$$(a|b) \equiv \sum_{j=0}^{\frac{N}{2}-1} 4\,\mathrm{Re}\left(\frac{\tilde{a}_j^* b_j}{S_n(f_j)}\Delta f\right) \tag{6}$$

where Δf is the frequency resolution, N is the number of samples, and S_n is the power spectral density (PSD) of the data segment.

There are cases where the assumption of stationary, Gaussian noise may not be valid due to the presence of glitches in data. These glitches may create a bias in the likelihood calculations, especially when they occur close to a gravitational-wave signal [6]. Therefore, segments of strongly non-stationary, non-Gaussian data

must be dealt with before the likelihood can be calculated using Eq. 5. While other methods have been developed to address non-stationarity over long periods of time [7–10], these do not account for glitches. It is also possible to model glitches and then remove them from data [11–19], but this can be time-consuming and is not guaranteed to fully remove all glitch impacts.

The increasing sensitivity of gravitational-wave detectors will increase the probability of glitches occurring near signals in upcoming observing runs. In this paper we present the implementation of inpainting as a method to mitigate the effect of glitches in PE analyses, specifically in `Bilby`. We find that, although this process significantly increases `Bilby`'s runtime, inpainting successfully prevents data inside a given segment from contributing to posteriors. Inpainting provides an alternative to more computationally expensive methods, and can be done within an analysis rather than as preprocessing step.

II. METHODS

The process of inpainting was derived in [8] and discussed in the context of PE in [20] and [21]. Here we review this derivation and its application to likelihood calculations.

We denote N_d the total number of data samples and N_h the number of data samples in the specified region to be inpainted, known as the *hole*. The hole is presumably within the chosen segment of data, such that $N_h \leq N_d$. The inpainting filter F is designed to satisfy

$$u^{(\alpha)T} C^{-1} F d = 0 \tag{7}$$

Here, $u^{(\alpha)}; \alpha = 0, ..., N_h - 1$ is a column vector of length N_d that equals 1 at one of the samples in the hole, and 0 elsewhere. Both the inverse noise covariance matrix C^{-1} and F are of size $N_d \times N_d$.

The covariance matrix is the Fourier transform of the PSD. Therefore, if the samples inside the hole satisfy Eq. 7 in the time domain, then these samples will not contribute to the noise-weighted inner products (Eq. 6) in the frequency domain.

To prevent the last term in Eq. 5 from creating a bias in the results, we must inpaint both the data and the waveform. For clarity, the term "waveform" will be used for waveforms generated within the analysis by sampling the parameter space, and the term "signal" will be used to refer to the gravitational wave signal that is presumably found in the data. By inpainting both the data and waveforms, the likelihood becomes

$$\mathcal{L}(d_{inp}|\theta) \propto \exp[-\frac{1}{2}(d_{inp}|d_{inp}) + (d_{inp}|h_{inp}) \tag{8}$$

$$-\frac{1}{2}(h_{inp}|h_{inp})] \tag{9}$$

where none of the terms have contributions from values inside the hole. Note that to ensure no contributions from inside of the hole, both the data and the waveform must be inpainted. With the likelihood completely blinded to the region inside the hole, the uncertainty in the results should increase but any glitch inside the hole will be prevented from creating a bias. Because this method is not dependent on the presence of glitches, applying an inpainting filter to stationary, Gaussian data should solely increase the uncertainty in the results.

We use two methods to calculate the inpainted data Fd, each of which has its own advantages and disadvantages.

A. The F method

The first method we use to obtain the inpainted data, which will be referred to as the "F method", involves directly computing the F matrix and applying it to the data. The F matrix is defined as

$$F \equiv I - AM^{-1}A^T C^{-1} \tag{10}$$

where $A_{i,\alpha} = u_i^{(\alpha)}$, and M is the $N_h \times N_h$ portion of the inverse covariance matrix corresponding to the samples to be inpainted. The resulting F matrix takes a form similar to an identity matrix with the exception that the rows corresponding to the samples in the hole have been altered.

The computation of F is done as a preprocessing step, and therefore does not contribute significantly to the run time of an analysis. The increase in run time is due to the matrix multiplication that inpaints the waveforms generated within `Bilby`. This matrix multiplication is of $\mathcal{O}(N_d^2)$, and occurs for every waveform that is sampled in the analysis.

An important factor to consider with this method is that a large amount of computer memory is required to store one or more $N_d \times N_d$ matrices. This is especially important when analyzing long signals such as neutron star black hole (NSBH) and binary neutron star (BNS) mergers. Because of this and the expected increase in run time for longer signals, we only use this method when analyzing BBH mergers. Because the full F matrix contains mostly zeroes and ones, this issue could be addressed by storing only the relevant rows of the F matrix (those corresponding to the samples inside the hole). However, we do not do this in this study.

B. The Toeplitz Method

From Eq. 10, the inpainting filter can also be interpreted as a filter that removes the projection of the data into the overwhitened data space inside the hole, leaving the values outside of the hole untouched. We can then rewrite our inpainted data as

$$Fd = d - d_{proj} \tag{11}$$

where d_{proj} is this projection. With Eq. 7

$$A^T C^{-1} d_{proj} = A^T C^{-1} d \tag{12}$$

We take d_{proj} to be zero everywhere except inside the holes. We can pick out the values in the hole by writing

$$A^T C^{-1} A A^T d_{proj} = A^T C^{-1} d \tag{13}$$

If there is only one hole to be inpainted, such that $A_{i,j} = A_{(i+1),(j+1)} = 1$ for some range of consecutive indices i, j and is 0 elsewhere, then $A^T C^{-1} A$ will be Toeplitz, and Eq. 13 can be solved for $A^T d_{proj}$ efficiently using Levinson recursion [22], which is done in $\mathcal{O}(N_h^2)$.

Algorithms for both of these inpainting methods were implemented in `Bilby` and their effects on run time and PE results were studied for a variety of cases.

III. VALIDATION TESTS

A. Time series tests

We first ensure that our algorithms perform as they were designed to by using them to inpainting segments of data and overwhitening the results. The overwhitened data should be zeroed inside of specified hole regions.

Fig. 1 shows a 4s segment of randomly generated data inpainted at three holes with different window lengths using the F method. The overwhitened strain has been effectively zeroed inside each hole.

Similarly, Fig 2 shows data that was inpainted using the Toeplitz method. This method can only be used for a single hole. The overwhitened strain is once again zeroed within the specified region. This confirms that the algorithms worked correctly.

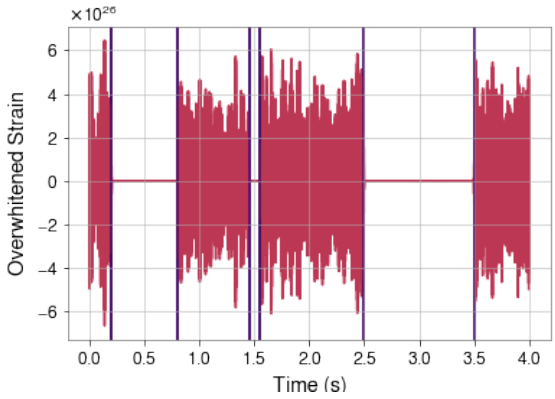

FIG. 1. Overwhitened inpainted data obtained using F method. Vertical lines show specified limits of each hole.

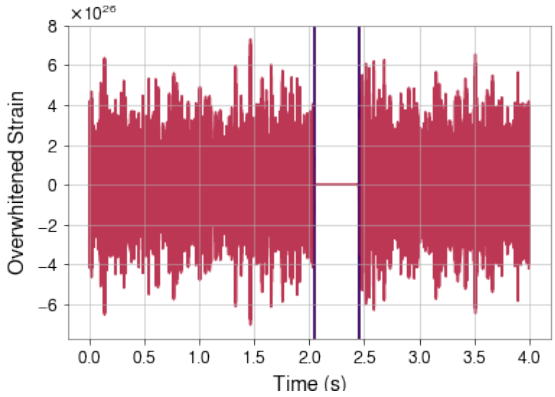

FIG. 2. Overwhitened inpainted data obtained using Toeplitz method. Vertical lines show the specified limits of the hole.

B. Efficiency

To study how each method increases `Bilby`'s run time, we measure the factor by which the single-likelihood evaluation time increases when the waveforms are inpainted. Using a 4 second segment of data sampled at 4096 Hz with an injected binary black hole (BBH) signal, the single-likelihood evaluation times of 30 samples are measured for 15 different window lengths between 0 and 2 seconds. We then take the median evaluation time for each window and divide by the median evaluation time for samples that were not inpainted. These factors of increase in single-likelihood evaluation times are shown in Fig. 3 for both methods.

The sharp spikes in this figure are most likely due to the small amount of samples that were evaluated. This could be improved by rerunning this test with a larger number of samples per window and a larger number of windows, although this would take a significant amount of time to complete. Nevertheless, this figure is enough to study the behavior of these algorithms.

Inpainting with the F method increases the evaluation time to over 100 times its original value for this segment duration and sample rate. As the inpainted

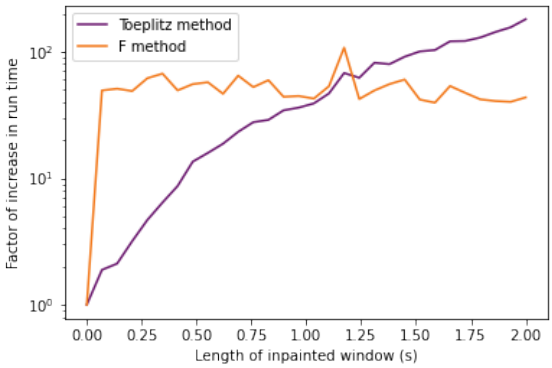

FIG. 3. Factor of increase in single-likelihood evaluation times (T/T_0) for different window lengths inpainted the using F method and the Toeplitz method.

window increases in length, the factor of increase stays approximately constant. This is because the matrix multiplication depends on N_d rather than N_h. On the other hand, because the Toeplitz method is in the order of N_h^2, the evaluation times are, for most windows, less than that of the F method, and increase as the window length increases.

Between a $1.00s$ and $1.25s$ window, the increase in evaluation time for the Toeplitz method surpasses that of the F method. This may be because the Toeplitz method has non-leading-order terms that are dependent on N_h, while the F method does not. Past this window length, these non-leading order terms may become more significant.

Based on these results, the Toeplitz method is preferred if there is only one hole to be inpainted and the window is not too wide. This is not only because it is faster, but also because less memory space is required. If there is more than one hole to be inpainted, then the F method should be used.

C. Injection Study

In order to confirm that inpainting stationary Gaussian data only increases uncertainty in the posteriors, we inject a binary black hole signal into Gaussian noise scaled to the PSD of both LIGO detectors at approximately 128 seconds before GW150914 [23]. We use the IMRPhenmoPv2 [24] approximant and 4 seconds of data sampled at 4096 Hz. The injected parameters for this signal are shown in Table I, along with the priors used. Spectrograms of the whitened data with the injection are shown in Fig. 4. The `Dynesty` sampler [25] was used to sample the parameter space.

TABLE I. Injection parameters and priors used.

Parameter	Value	Prior Shape	Limits	Boundary
M_{chirp}	15.53	Uniform	10-20	-
q	0.52	Uniform	0.125-1	-
a_1	0.65	Uniform	0-0.99	-
a_2	0.65	Uniform	0-0.99	-
ϕ_{12}	0.0	Uniform	0-2π	Periodic
ϕ_{JL}	0.0	Uniform	0-2π	Periodic
d_L	100	Uniform	50-2000	-
δ	1.00	Cosinusoidal	-	-
α	2.00	Uniform	0-2π	Periodic
θ_{JN}	1.65	Sinusoidal	-	-
ψ	1.50	Uniform	0-π	Periodic
ϕ	2.00	Uniform	0-2π	Periodic
t_c	2.5	Uniform	2.4-2.6	-
θ_1	0.0	Sinusoidal	-	-
θ_2	0.0	Sinusoidal	-	-

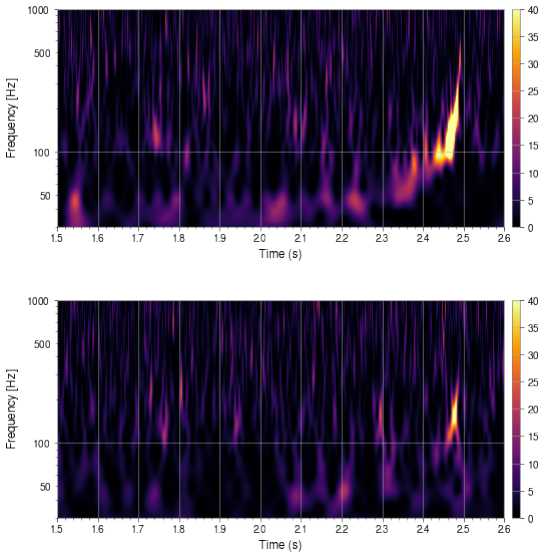

FIG. 4. Whitened spectrograms of injected BBH signal in LIGO-Hanford (top) and LIGO-Livingston (bottom).

1. *Parameter recovery*

To provide a baseline for the results, a standard analysis without the inpainting filter was performed. A full PE analysis was done inpainting the data of both detectors at a region centered 0.1s before the injected geocentric time with a 0.05s window. The whitened spectrograms of the inpainted data with this configuration is shown in Fig. 5. A second inpainted analysis was done centered 0.25 seconds before the injected geocentric time with a 0.2 second window. The whitened spectrograms of the inpainted signal with this configuration are shown in Fig. 6. Because we only chose one

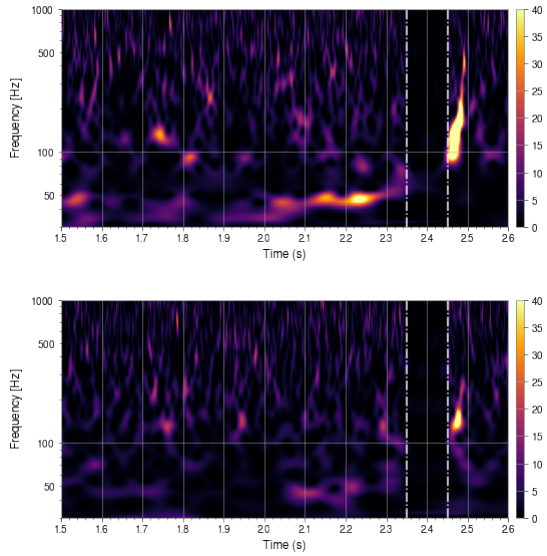

FIG. 5. Whitened spectrograms of injected BBH signal in LIGO-Hanford (top) and LIGO-Livingston (bottom), inpainted 0.1 s before t_c with a 0.05 s window.

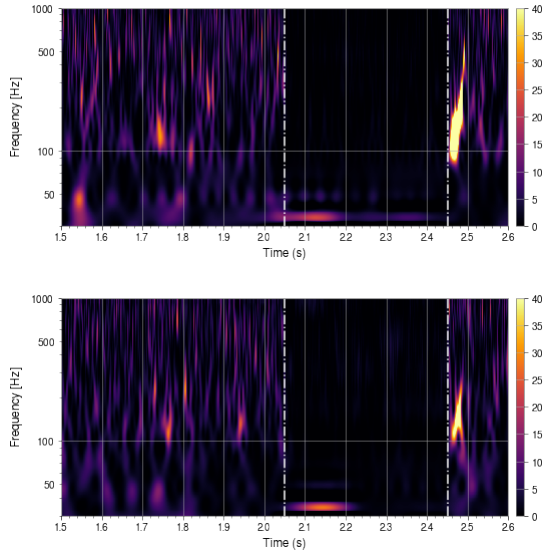

FIG. 6. Whitened spectrograms of injected BBH signal in LIGO-Hanford (top) and LIGO-Livingston (bottom), inpainted 0.25 s before t_c with a 0.2 s window.

window for each inpainted analysis and both window lengths are short, both inpainted analyses were done with the Toeplitz method.

TABLE II. Recovered parameters for the standard analysis and the inpainted analyses with both configurations.

	Injected	Standard (w=0)	Inpainted (w=0.01)	Inpainted (w=0.2)
$\mathcal{M}/M_\odot$	15.53	$15.63^{+0.12}_{-0.11}$	$15.56^{+0.15}_{-0.14}$	$15.77^{+0.73}_{-0.43}$
q	0.52	$0.60^{+0.16}_{-0.12}$	$0.55^{+0.11}_{-0.09}$	$0.53^{+0.26}_{-0.13}$
a_1	0.65	$0.81^{+0.11}_{-0.15}$	$0.90^{+0.06}_{-0.11}$	$0.84^{+0.09}_{-0.13}$
a_2	0.65	$0.55^{+0.24}_{-0.24}$	$0.34^{+0.26}_{-0.23}$	$0.58^{+0.26}_{-0.29}$
ϕ_{12}	0.0	$3.25^{+2.69}_{-2.89}$	$3.07^{+2.88}_{-2.78}$	$3.10^{+2.11}_{-2.08}$
ϕ_{JL}	0.0	$3.74^{+0.45}_{-0.47}$	$3.16^{+2.15}_{-2.11}$	$3.60^{+2.15}_{-3.23}$
d_L/Mpc	100	$104.41^{+15.47}_{-13.49}$	$104.49^{+17.54}_{-14.42}$	$89.09^{+16.24}_{-14.13}$
δ	1.00	$-0.93^{+0.08}_{-0.07}$	$-0.93^{+1.01}_{-0.08}$	$0.99^{+0.11}_{-0.86}$
α	2.00	$4.98^{+0.18}_{-0.25}$	$4.89^{+0.29}_{-0.46}$	$2.02^{+3.78}_{-0.32}$
θ_{JN}	1.65	$1.45^{+0.07}_{-0.07}$	$1.48^{+0.09}_{-0.10}$	$1.56^{+0.09}_{-0.13}$
ψ	1.50	$1.66^{+0.12}_{-1.49}$	$0.28^{+2.10}_{-0.18}$	$0.48^{+2.51}_{-0.41}$
ϕ	2.00	$2.87^{+2.48}_{-2.04}$	$3.10^{+2.22}_{-2.15}$	$3.10^{+2.20}_{-2.10}$
t_{geo}/s	2.5	$2.48^{+0.00}_{-0.00}$	$2.48^{+0.02}_{-0.00}$	$2.50^{+0.00}_{-0.00}$
θ_1	0.0	$0.26^{+0.13}_{-0.11}$	$0.24^{+0.14}_{-0.10}$	$0.26^{+0.20}_{-0.14}$
θ_2	0.0	$0.67^{+0.37}_{-0.31}$	$1.07^{+0.65}_{-0.57}$	$0.70^{+0.45}_{-0.33}$

2. *Comparison to standard posterior*

The recovered parameters from the standard and inpainted runs, along with their 90% credible intervals are shown in Table II. In the standard analysis, most parameters were recovered within the 90% credible interval. However, parameters such as sky location, geocentric time, and angles related to angular momenta and spin were not recovered within this credible interval. The error in sky location may be because time delay is a more reliable measurement, and `Bilby` recovers a sky location within the ring of sky locations with equal time delays. Meanwhile, ϕ_{12}, ϕ_{JL}, θ_{JN}, θ_1, and θ_2 are likely not in their 90% credible intervals due to the priors used for these parameters. To obtain better recovery of parameters, these analyses should be done for a set of injections with parameters that are representative of the priors used.

Although the accurate recovery of parameters is important for both standard and inpainted analyses, once we establish that the posteriors obtained in the standard analysis are reasonable, we are more interested in how the posteriors obtained with the inpainted analysis compare with those obtained with the standard analysis. The second and third columns of Table II show the results for both inpainted analyses.

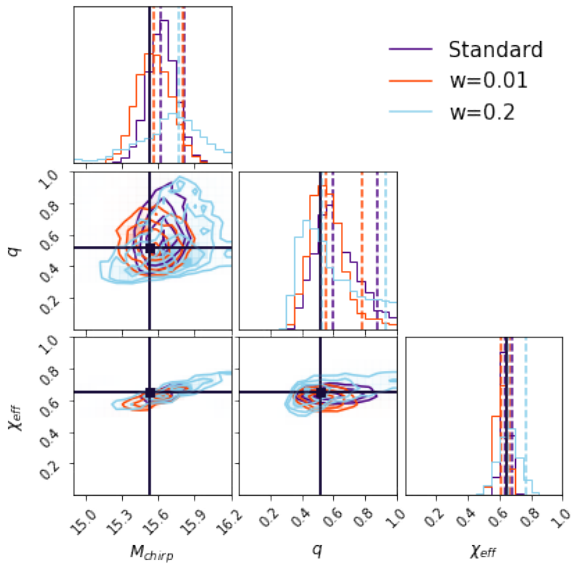

FIG. 7. Posteriors for chirp mass, mass ratio, and effective spin for the standard, first inpainted analysis (window=0.01s), and second inpainted analysis (window=0.2s). Injected values are shown in black.

The corner plots for the chirp mass, mass ratio, and effective spin posteriors are shown in Fig. 7. These parameters are typically recovered with the most accuracy, so it is important to verify that inpainting does not introduce a bias when recovering them. As seen in the figure, the corner plots obtained with the first inpainted configuration are similar to those obtained with the standard analysis. This reveals that the implementation of inpainting did modify the data, but the window was too small to create large differences in the posterior, which was the expected result for a window size of 0.01 seconds. Once the window was increased to 0.2 seconds, the posteriors appear to spread out in relation to the standard results. The median of these posteriors did not significantly shift in a way that would suggest a bias, which confirms that inpainting only increased the uncertainty in the results.

For further comparison and validation of our methods, Table III shows the statistical properties of the log-likelihood posteriors obtained for each of the three analyses. These results suggest that the log-likelihood decreases as the inpainted window increases. This is because when larger segments of the data are inpainted, there are more terms that are zeroed out in the likelihood calculation, decreasing its final value. For the same reason, the standard deviation increases as the window increases. Data that would originally make the uncertainty smaller is being removed with the filter.

TABLE III. Statistical properties of log-likelihood posteriors.

Quantity	Standard (w=0s)	Inpainted (w=0.01s)	Inpainted (w=0.2s)
Median	246.600	196.433	169.047
Standard Deviation	2.451	2.651	2.935

3. *Effect of changing center time*

As shown in [6], glitches have a particularly strong effect on PE results when they occur at or near merger time. For this reason, we study the effect of inpainting as the center time of the hole approaches merger time.

Fig. 8 shows chirp mass posteriors obtained by inpainting a 0.05 s window for different time delays, from 0 seconds (directly over merger time) to 0.5 seconds with a step of 0.1 seconds. Delay refers to the amount of time before the merger at which the center of the hole was placed. The gray distribution shows the results for the standard analysis. It is important to note that for these results, all parameters except chirp mass, mass ratio, phase, and geocentric time were fixed to reduce run time.

The posteriors obtained for delays between 0.5 and 0.1 seconds appear comparable to the standard result. However, the posterior obtained when inpainting directly over the merger time is significantly different from the rest. The abrupt change between the delay of 0.1 and 0 may be, in part, because not enough delays were tested between these values to produce a smooth transition. BBH signals are short and therefore the 0.1 second step used may have been too large in relation to the signal duration. Additionally, fixing various parameters may have caused the analyses at lower delays to better recover the chirp mass even when the inpainted window was close to the merger time.

The effect of changing the center time should be restudied by sampling the entire parameter space and by including more delay values. This will better determine whether the flattened posterior obtained when inpainting over merger time is due to the posterior approaching the prior distribution, which is uniform for chirp mass, or to a bias in the PE analysis.

4. *Effect of changing window length*

In Sec. III C 2 the results suggested that as the window length increases, the posteriors appear to spread out in relation to the standard results. To further look into this behavior, shorter analyses were done with

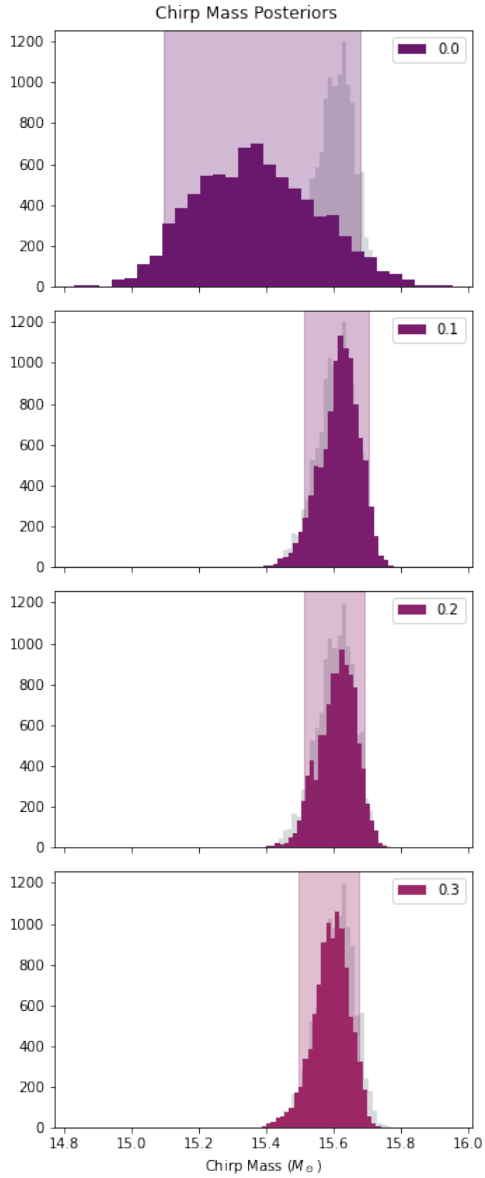

FIG. 8. Posteriors obtained for chirp mass by inpainting a 0.05s window centered 0.0s to 0.3s before merger time.

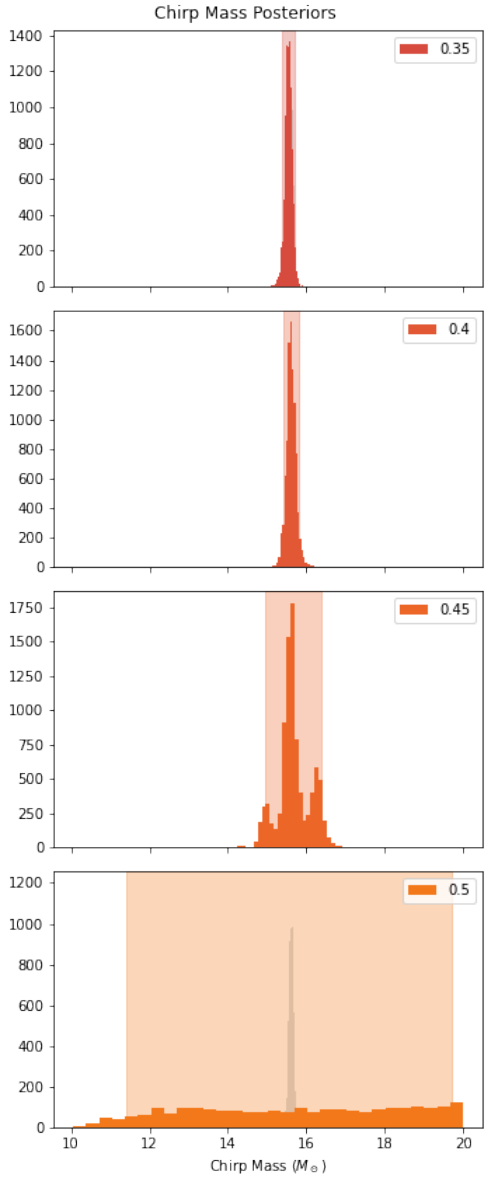

FIG. 9. Posteriors obtained for chirp mass by inpainting a hole centered 0.5s before merger time with windows ranging from 0.35s to 0.5s.

changing window lengths, with the idea that inpainting the entire data segment should result in posterior distributions that resemble the priors used.

Fig. 9 once again shows chirp mass posteriors compared to the standard posterior in gray for analyses with different window lengths, ranging from 0.35s to 0.5s. These holes were centered 0.5 seconds before merger time. The results for data inpainted with a window of 0.35s were in agreement with those of the standard analysis. However, as we increase the inpainted window, these posteriors flatten. As expected, past a certain window length the posteriors aligned with the prior used, which was uniform in chirp mass and limited from 10 $M_\odot$ to 20 $M_\odot$.

For this particular data segment with the center time placed 0.5s before the geocentric time of the merger, the threshold window length with which the analysis is still able to recover meaningful posteriors lies between 0.45s and 0.5s. This threshold's depedency on waveform types and on center time of the hole can be further studied.

5. *Likelihood Reweighting*

The significant increase in run time explained in Sec. III B can be improved by a reweighting procedure as explained in [26]. When this procedure is used for inpainting, an initial analysis can be performed only inpainting the data, and not the waveforms, such that the likelihood,

$$\mathcal{L}_{\oslash}(d_{inp}|\theta) \propto \exp[-\frac{1}{2}(d_{inp}|d_{inp}) + (d_{inp}|h) - \frac{1}{2}(h|h)] \tag{14}$$

is allowed to have a bias attributed to the last term. Assuming that this bias is not too large, the resulting posterior samples can be reweighted using a weight function

$$w(d|\theta) \equiv \frac{\mathcal{L}(d_{inp}|\theta)}{\mathcal{L}_{\oslash}(d_{inp}|\theta)} \tag{15}$$

such that

$$p(\theta|d) = w(d|\theta)\frac{\mathcal{L}_{\oslash}(d|\theta)\pi(\theta)}{z(d)} \tag{16}$$

This reweighting procedure has already been implemented in `Bilby`, and it is straightforward to apply it with inpainting. To test how a posterior improves with the rewighting procedure, we inpaint our Livingston data as a pre-processing step and then run a standard analysis to obtain initial likelihoods in the form of Eq. 14. Then we reweight these results with an inpainted analysis using the built-in rewighting function. These results can be compared to the original fully inpainted posteriors to observe whether reweighting improves the biased results.

Fig. 10 shows these results. A close look at this figure shows that the uncertainty region of the reweighted analysis is closer to that of the fully inpainted analysis than that of the biased analysis, which was the desired effect of the reweighting procedure.

Although much more efficient than a fully inpainted analysis, reweighting must be done with caution. If the bias in the likelihood in Eq. 14 is too large, then the resulting posterior will not have enough samples within the true credible regions to be reighted correctly. This can happen if the window is too large in relation to the signal.

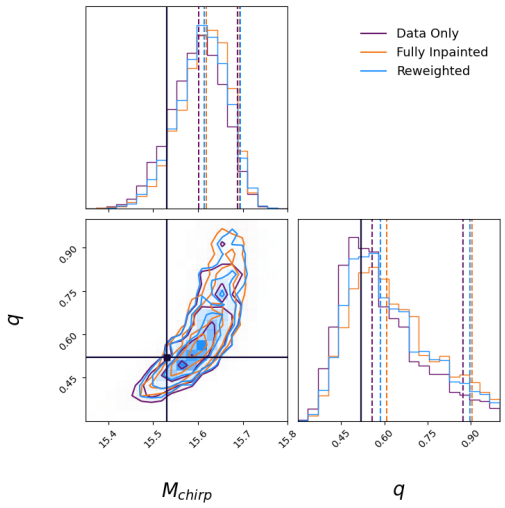

FIG. 10. Posteriors obtained for chirp mass and mass ratio by inpainting only the data as a preprocessing step (*Data Only*), inpainting both the data and the waveforms in the analysis (*Fully Inpainted*) and reweighting the posterior samples from the analysis with the bias (*Reweighted*).

IV. FUTURE WORK

With the results that have been discussed, there is a foundation set for further study of the applications of inpainting. The current algorithms can be officially added onto `Bilby` for testing with `Bilby PIPE`, which will allow for more thorough analyses and conclusions.

Because PE is already an extensive process, it is important to improve the efficiency of our algorithms wherever it is possible. As previously mentioned, the F method can be improved by rewriting the algorithm such that it stores only the rows of the F matrix corresponding to the samples inside the hole. Once this is implemented, this method should be of $\mathcal{O}(N_d N_h)$ instead of $\mathcal{O}(N_d^2)$, which makes it much more usable for future analyses.

More inpainted analyses should also be performed on a population of injections that is representative of the priors used. This will provide a better measure of how often inpainted analyses recover injected values. The variables that can be tested include the center time and window of the holes, the amount of holes, the distance between holes, and the waveform types to be analyzed. With this study, we can be more certain that this method does not create a bias in a PE analysis.

The purpose of inpainting is to mitigate the effect of glitches, so future analyses should be done with injected glitches to study if and how the posteriors improve when a glitch is inpainted. This will also help determine how much of the glitch must be inpainted to obtain better results. Once this is understood, inpainting can be applied to detected signals that have occurred near glitches, such as GW170817, and will be ready for use in O4 and future observing runs.

V. ACKNOWLEDGEMENTS

This project would not have been possible without the guidance and support of my mentor, Derek Davis, for whom I am immensely thankful. I also thank the LIGO group at the California Institute of Technology for their input on this project and for the positive work environment that they provided. I gratefully acknowledge the National Science Foundation, the Student-Faculty Programs office, the LIGO Summer Undergraduate Research Fellowship, the Victor M. Blanco Fellowship, and my home institution, the University of Puerto Rico at Mayagüez, for making this experience possible.

[1] D. Davis and M. Walker, Detector Characterization and Mitigation of Noise in Ground-Based Gravitational-Wave Interferometers, Galaxies **10**, 12 (2022).

[2] G. Ashton *et al.*, BILBY: A user-friendly Bayesian inference library for gravitational-wave astronomy, Astrophys. J. Suppl. **241**, 27 (2019), arXiv:1811.02042 [astro-ph.IM].

[3] I. M. Romero-Shaw *et al.*, Bayesian inference for compact binary coalescences with bilby: validation and application to the first LIGO–Virgo gravitational-wave transient catalogue, Mon. Not. Roy. Astron. Soc. **499**, 3295 (2020), arXiv:2006.00714 [astro-ph.IM].

[4] B. P. Abbott *et al.* (LIGO Scientific, Virgo), A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, Class. Quant. Grav. **37**, 055002 (2020), arXiv:1908.11170 [gr-qc].

[5] P. Whittle, *Hypothesis testing in time series analysis*, Vol. 4 (Almqvist & Wiksells boktr., 1951).

[6] J. Powell, Parameter Estimation and Model Selection of Gravitational Wave Signals Contaminated by Transient Detector Noise Glitches, Class. Quant. Grav. **35**, 155017 (2018), arXiv:1803.11346 [astro-ph.IM].

[7] O. Edy, A. Lundgren, and L. K. Nuttall, Issues of mismodeling gravitational-wave data for parameter estimation, Phys. Rev. D **103**, 124061 (2021), arXiv:2101.07743 [astro-ph.IM].

[8] B. Zackay, T. Venumadhav, J. Roulet, L. Dai, and M. Zaldarriaga, Detecting gravitational waves in data with non-stationary and non-Gaussian noise, Phys. Rev. D **104**, 063034 (2021), arXiv:1908.05644 [astro-ph.IM].

[9] S. Mozzon, L. K. Nuttall, A. Lundgren, T. Dent, S. Kumar, and A. H. Nitz, Dynamic Normalization for Compact Binary Coalescence Searches in Non-Stationary Noise, Class. Quant. Grav. **37**, 215014 (2020), arXiv:2002.09407 [astro-ph.IM].

[10] S. Kumar, A. H. Nitz, and X. J. Forteza, Parameter estimation with non stationary noise in gravitational waves data, (2022), arXiv:2202.12762 [astro-ph.IM].

[11] N. J. Cornish and T. B. Littenberg, BayesWave: Bayesian Inference for Gravitational Wave Bursts and Instrument Glitches, Class. Quant. Grav. **32**, 135012 (2015), arXiv:1410.3835 [gr-qc].

[12] C. Pankow *et al.*, Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817, Phys. Rev. D **98**, 084016 (2018), arXiv:1808.03619 [gr-qc].

[13] N. J. Cornish, T. B. Littenberg, B. Bécsy, K. Chatziioannou, J. A. Clark, S. Ghonge, and M. Millhouse, BayesWave analysis pipeline in the era of gravitational wave observations, Phys. Rev. D **103**, 044006 (2021), arXiv:2011.09494 [gr-qc].

[14] K. Chatziioannou, N. Cornish, M. Wijngaarden, and T. B. Littenberg, Modeling compact binary signals and instrumental glitches in gravitational wave data, Phys. Rev. D **103**, 044013 (2021), arXiv:2101.01200 [gr-qc].

[15] S. Hourihane, K. Chatziioannou, M. Wijngaarden, D. Davis, T. Littenberg, and N. Cornish, Accurate modeling and mitigation of overlapping signals and glitches in gravitational-wave data, Phys. Rev. D **106**, 042006 (2022), arXiv:2205.13580 [gr-qc].

[16] J. Merritt, B. Farr, R. Hur, B. Edelman, and Z. Doctor, Transient glitch mitigation in Advanced LIGO data, Phys. Rev. D **104**, 102004 (2021), arXiv:2108.12044 [gr-qc].

[17] W. Wei and E. A. Huerta, Gravitational Wave Denoising of Binary Black Hole Mergers with Deep Learning, Phys. Lett. B **800**, 135081 (2020), arXiv:1901.00869 [gr-qc].

[18] A. Torres-Forné, E. Cuoco, J. A. Font, and A. Marquina, Application of dictionary learning to denoise LIGO's blip noise transients, Phys. Rev. D **102**, 023011 (2020), arXiv:2002.11668 [gr-qc].

[19] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. Di Renzo, and G. Ashton, Subtracting glitches from gravitational-wave detector data during the third observing run, (2022), arXiv:2207.03429 [astro-ph.IM].

[20] J. Y. L. Kwok, R. K. L. Lo, A. J. Weinstein, and T. G. F. Li, Investigation of the effects of non-Gaussian noise transients and their mitigation in parameterized gravitational-wave tests of general relativity, Phys. Rev. D **105**, 024066 (2022), arXiv:2109.07642 [gr-qc].

[21] C. D. Capano, M. Cabero, J. Westerweck, J. Abedi, S. Kastha, A. H. Nitz, A. B. Nielsen, and B. Krishnan, Observation of a multimode quasi-normal spectrum from a perturbed black hole, (2021), arXiv:2105.05238 [gr-qc].

[22] N. Levinson, The wiener (root mean square) error criterion in filter design and prediction, Journal of Mathematics and Physics **25**, 261 (1946), https://onlinelibrary.wiley.com/doi/pdf/10.1002/sapm1946251261.

[23] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.*, Observation of gravitational

waves from a binary black hole merger, Physical review letters **116**, 061102 (2016).

[24] M. Hannam, P. Schmidt, A. Bohé, L. Haegel, S. Husa, F. Ohme, G. Pratten, and M. Pürrer, Simple model of complete precessing black-hole-binary gravitational waveforms, Physical review letters **113**, 151101 (2014).

[25] J. S. Speagle, dynesty: a dynamic nested sampling package for estimating bayesian posteriors and evidences, Monthly Notices of the Royal Astronomical Society **493**, 3132 (2020).

[26] E. Payne, C. Talbot, and E. Thrane, Higher order gravitational-wave modes with likelihood reweighting, Phys. Rev. D **100**, 123017 (2019), arXiv:1905.05477 [astro-ph.IM].