Exploring the Impact of Frequency Cuts on Gravitational-Wave Parameter Estimation

Sena Kalabalık
Boğaziçi University, Turkey

Mentors: Lucy M. Thomas, Rhiannon Udall and Derek Davis
LIGO, California Institute of Technology, Pasadena

Interim Report I
LIGO Caltech SURF Program 2024

Gravitational-wave science has empowered a new era in astrophysical exploration, with the characterization of gravitational-wave signals holding crucial importance. However, accurately characterizing these signals poses challenges, especially in the presence of short-duration noise transients, known as "glitches." A specific strategy employed to address this issue and mitigate the effects of glitches involves cutting a particular frequency range off the signal. In this study, we aim to delve into the critical role that frequency cuts play in gravitational-wave analysis. Through systematic analysis and simulation-based experiments, incorporating data injections and utilizing a neural posterior estimator, we will investigate the effects of different frequency cut configurations on our interpretations of parameter posterior distributions across a diverse array of characteristic waves.

I. INTRODUCTION

Gravitational waves originate predominantly from the accelerated motion of massive objects, such as the orbital movement of black holes and neutron stars. This motion disturbs the fabric of space-time, leading to the propagation of waves in all directions from the source. Traveling at the speed of light, these cosmic ripples convey valuable information about their source's characteristics and provide insights into the fundamental nature of gravity. As gravitational waves travel through space, they induce tiny expansions and contractions in the spatial dimensions they traverse. The measurements of this strain are made possible by the Laser Interferometer Gravitational-Wave Observatory (LIGO) and the Virgo Collaboration (LVC). The LVC operates three detectors: two LIGO detectors located in the United States and one Virgo detector in Italy. The detectors are specialized versions of a Michelson interferometer, which simply consists of two equally long, perpendicular arms with mirrors at their ends. A laser beam is split and sent down each arm, reflecting off the mirrors and then recombining at a central detector. When a gravitational wave passes through, it causes the lengths of the arms to change slightly. This alters the interference pattern when the laser beams recombine, allowing scientists to detect the gravitational waves in the form of weak signals.

One of the major sources of gravitational waves are binary black hole mergers and neutron star mergers, which fall under the category of compact binary coalescences (CBC).The characterization of these waves is especially crucial since they carry promising information about the nature of compact objects. When a signal is detected, the source characterization is made possible by employing Bayesian inference, which relies on having established models for both the signals and the detector noise. In the context of gravitational waves, these signal models are represented by waveform predictions h(θ), which are contingent upon various source parameters θ, such as masses and locations of the objects. Meanwhile, the detector noise is typically assumed to be stationary and Gaussian, characterized by a certain spectrum that can be empirically estimated. Collectively, these models yield the likelihood $p(d|\theta)$ for the observed strain data d, presumed to comprise both signal and noise components. By selecting a prior $p(\theta)$ over the parameters, the posterior distribution is determined through Bayes' theorem as:

$$p(\theta|d) = \frac{p(d|\theta)p(\theta)}{p(d)}, \tag{1}$$

where $p(d)$ acts as a normalizing factor termed the evidence. This posterior distribution encapsulates our beliefs regarding the source parameters, given the observed data.[1]

In order to accomplish the inference, we have chosen to utilize the DINGO method [1] over other conventional methods commonly used in LVC, such as LALInference [2] and Bilby [3]. The primary reason for this choice is the imperative need for speed in conducting numerous trials of parameter estimations. Traditional methods such as LALInference and Bilby employ stochastic algorithms like Markov chain Monte Carlo (MCMC) to characterize the posterior distribution by drawing samples from it. However, these algorithms are computationally intensive, demanding numerous likelihood evaluations for each independent posterior sample. Each likelihood evaluation necessitates a waveform simulation, making the entire inference process time-consuming. However, DINGO is a neural posterior estimation (NDE) model, and has shown to achieve both significantly reduced analysis time and high accuracy. As a likelihood-free and simulation-based inference model, the simple principle of DINGO is to generate numerous simulated datasets, each with its corresponding parameters, and utilize these datasets to train a specific type of neural network called a normalizing flow, which is employed to approximate the posterior distribution. Once the network is trained, it can rapidly produce new posterior samples following a detection.[1]

Accurate source characterization depends on the specific assumptions about the behavior of the detector noise, however, these assumptions are violated when there are short-duration noise transients, called "glitches" present in the signal [4]. Glitches can arise from various reasons such as instrumental artifacts or environmental disturbances, however, their sources are not typically directly identifiable, therefore it is not uncommon for glitches to occur unpredictably in the signal. Throughout their third observing run (O3) [5, 6], the median rate of glitches in the LIGO and Virgo detectors has been reported to have surpassed 1 per minute for the majority of the duration, which indicates that the coincidence of the glitches with gravitational-wave signals are expected quite commonly [7]. Since these glitches corrupt the signal and invalidate the noise assumptions for typical gravitational-wave source characterization processes, their identification and mitigation are crucial for accurate analysis.

The process of identifying and subtracting glitches is not straightforward, leading to the adoption of various strategies. One of the most common and sophisticated methods is the BayesWave algorithm [8], which models the glitch and subtracts it by only using the strain data. The algorithm assumes that the strain data consists of Gaussian noise, a gravitational-wave signal, and a glitch. It models the glitch as a sum of sine-Gaussian wavelets, and these wavelets are marginalized over the parameters using a trans-dimensional Markov chain Monte Carlo (MCMC). As the result of the algorithm, a posterior distribution of time series of the glitch is obtained, and the mitigation is conducted by randomly selecting a sample from the posterior, and subtracting it from the data [7]. Despite the common use of the BayesWave algorithm, various studies, including [9] and [5], have demonstrated that the method may potentially leave residual artifacts within the signal, which is not unexpected since the probability of a randomly drawn sample wavelet accurately capturing the precise characteristics of the glitch is low.

Another approach for glitch mitigation is gwsubstract algorithm [10], which uses information from auxiliary channels. The algorithm assumes that the measured strain is a linear combination of time series from different sources, where one of these sources can be modeled as the convolution of a witness time series and an unknown transfer function. In this approach, the transfer function between the auxiliary sensor and the strain data channel is determined and used to estimate the contribution of the noise source to the strain data. However, the accuracy of this subtraction method depends on the accuracy of the auxiliary sensor and the transfer function estimate [7, 9]. The mentioned systematic and statistical uncertainties of these two methods poses the risk that even after BayesWave or subtract algorithm is used, the data can still be undersubstracted, as shown in [5, 9].

Although the glitches are commonly dealt with by eliminating data in the time domain, it is possible to come across cases where the glitch only affects a specific frequency range. In such instances, addressing data quality issues might involve employing a narrower frequency range during the analysis [11]. This strategy has been adopted in the cases of [5, 9]. In [5], it has been reported that after the application of these methods, the identified glitch is considered unmitigated if the data surrounding the event are inconsistent with Gaussian noise. In such cases, they evaluated the SNR lost by restricting the frequency range of data considered in parameter inference to fully remove the glitch. If the SNR loss is below 10%, they used the reduced frequency range in the analyses. Otherwise, they have used the nominal frequency range. Similarly in [9], to further investigate the relation between the potentially under-substractive glitch mitigation strategies and the estimated spin-precession posteriors, they have limited the frequency range above a progressively increasing lower limit, and evaluated the SNR values as well as posteriors of various parameters for each lower limit. The results showed that even though SNR loss was small, the posterior for the spin-precession parameter χ_{P} became less informative when a more increased lower limit was used, possibly indicating that the glitch remnants were causing a misleading estimation of the posterior.

While previous studies, including [9] and [5], have explored the benefits of constraining the frequency range of data primarily for mitigating the remnants of glitches, our investigation suggests that frequency cuts may have overlooked implications. An example of this concern can be found in [7]. In this study, all three methods described have been employed and compared by their SNR values and posterior distributions. When the lower frequency limit was raised, the posteriors for spin parameters became less informative, and more influenced by the prior. Expectedly, the SNR loss was significant when higher limits were placed on the lower frequency. The results concluded that the glitch subtraction strategies narrowed the posterior distributions, and only caused a little change in SNR, therefore indicating their superiority over frequency cuts. However, of particular interest is the unexpected revelation that different frequency cut configurations led to distinctly different estimations for the same parameter, implying a complex scenario. The figures obtained demonstrated a significant differentiation in the posteriors for the χ_{eff} parameter, with noticeable shifts across the plot, when different frequency range limits were applied. This demonstration implies a sophisticated relation between frequency cuts and source characterization for gravitational-wave signals.

II. OBJECTIVES

The primary objective of this research is to analyze the impact of different frequency cuts on parameter estimation results for compact binary coalescences. This involves systematically examining how variations in frequency cuts influence the posterior distributions of gravitational wave signals. Given that current frequency cut strategies for noise mitigation lack a systematic approach, this study aims to establish a more educated and reliable framework, potentially reducing incorrect estimations in PE results.

Additionally, this research seeks to explore the feasibility and effectiveness of utilizing a deep learning (DL) model, DINGO, for the PE pipeline. By comparing the results obtained from the conventional PE pipeline, Bilby, with those derived from DINGO, the study aims to evaluate DINGO's reliability and efficiency for large-scale analysis. This investigation is not only significant for our analysis, but also for the ongoing observing run, O4, where traditional methods may be limited by the large amount of computations needed. A successful integration of DL algorithms could enhance the speed and the efficiency of parameter estimation, thereby improving overall analysis capabilities.

III. CURRENT PROGRESS

To achieve my research objectives, I first aimed to comprehend the Dingo pipeline to prepare for utilizing it in subsequent steps of my research. I analyzed the event GW150914 by following an official Dingo tutorial [12]. This analysis involved four steps: generating a waveform dataset, generating an ASD dataset, training the network, and performing inference. These steps required only configuration file edits, not custom code manipulation, which is advantageous for handling a large number of tasks in the more advanced applications.

In the first step, the configuration file included information and preferences for the generated waveforms, such as priors, waveform generator specifications, and domain specifications. For simplicity, no compression was applied. The configuration settings specified the frequency domain (`f_min`, `f_max`, `delta_f`), the waveform generator (approximant: `IMRPhenomD`, `f_ref`: 20.0), and intrinsic priors (e.g., mass constraints, spin components, and chirp mass).

The second step required setting another configuration file to generate the ASD dataset. The settings in `asd_dataset_settings.yaml` included attributes like sampling frequency (`f_s`), time segment length (`time_psd`), segment duration (`T`), and window type (`Tukey window`). These settings defined how the ASD was estimated using segments of noise data from detectors H1 and L1 during the O1 observing run. For computational efficiency in this simple model, only one ASD was generated and used for noise realization.

For network training, it was important to specify the neural network's features, which consist of two components: a normalizing flow (neural spline flow) and an embedding network. Key settings for the model included the number of flow transforms and hidden dimensions. The training configuration specified hyperparameters such as the number of epochs and optimizer settings. This approach considered extrinsic parameters during training to reduce computational costs, rather than generating waveforms for each parameter set.

In the final step, inference was performed using the `dingo_pipe` command, allowing for the analysis and manipulation of the results. The resulting posterior distributions for the 11 parameters of GW150914 are shown in Figure 1. While the results were not highly promising, this was expected due to the simplified workload for this analysis.

Another milestone in preparing for my analysis involved becoming proficient with Bilby and mastering essential tools and classes for data manipulation. I utilized the `prior` class to construct, evaluate, and sample from priors, while also creating insightful plots. Additionally, working with the `WaveformGenerator` and `Interferometer` classes enabled me to bridge theoretical waveforms with realistic detector observations, highlighting key detector features affecting waveform projections and noise realizations.

One of the pivotal concepts I explored was noise realization. This included loading the actual PSD of the event GW191109 into the detector simulation, using this ASD to accurately set the strain data and simulate realistic noise conditions. As part of these processes, I worked in both time domain (TD) and frequency domain (FD), honing my skills in performing effective Fourier and Inverse Fourier transforms using Bilby. I also employed tools such as whitening to enhance plot readability and interpretability.

Furthermore, I gained proficiency in injecting signals into detector simulations, illustrating these findings through plotted comparisons in both TD domain and spectrogram, as shown in Figure 2 and Figure 3 respectively. As part of my exercises, I also applied a strategic method of subtracting injected waveforms to analyze residuals. I plotted these in Figure 4 and Figure 5 to assess their compatibility and verify alignment with expected noise levels.

Weekly seminars have also been quite productive for my learning journey, providing invaluable insights into CBC parameter estimation methodologies, modeling and characterization of the detectors, solving equations numerically with Python, and testing general relativity (GR). These sessions have also enhanced my practical skills, such as navigating computational tools and efficiently running codes on both my local machine and the computer clusters.

Another enriching experience was our visit to the Hanford LIGO site, where I gained firsthand exposure to the intricate workings of gravitational wave detectors. This visit not only bolstered my enthusiasm for the field but also deepened my understanding of detector operations and the complexities involved in gravitational wave detection.

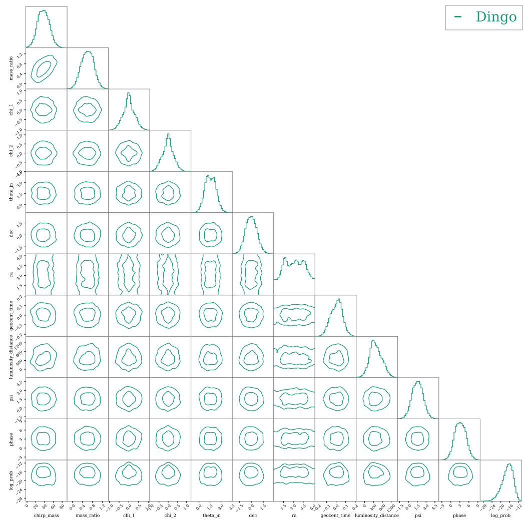

FIG. 1. This figure illustrates the posterior distribution results for 11 parameters of GW150914. The less promising results are anticipated due to the simplified workload for this analysis.

IV. NEXT STEPS

The planned next steps outlined for this project are as follows:

- Prepare 100 injections to be processed by the pipelines.
- Train two DINGO networks for each frequency cut: one using a single ASD for computational efficiency and the other using multiple ASDs for more robust results.
- Conduct inference with both DINGO and Bilby for each frequency cut range, and compare the results to assess DINGO's reliability in large-scale analysis.
- Continue inference with DINGO based on the results to derive posterior distributions for the injections at each frequency cut.
- Analyze the outcomes to develop systematic conclusions, improving comprehension and strategies for implementing frequency cuts.
- Aim to utilize DINGO for analyzing O4 data.

V. ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation Research Experience for Undergraduates (NSF REU) program, the LIGO Laboratory Summer Undergraduate Research Fellowship program (NSF LIGO), and the California Institute of Technology Student-Faculty Programs.

[1] Maximilian Dax, Stephen R. Green, Jonathan Gair, Jakob H. Macke, Alessandra Buonanno, and Bernhard Schölkopf. Real-time gravitational wave science with neural posterior estimation. *Phys. Rev. Lett.*, 127:241103, 2021.

[2] J. Veitch et al. Parameter estimation for compact binaries with ground-based gravitational-wave observations using the lalinference software library. *Phys. Rev. D*, 91:042003, 2015. Published 6 February 2015.

[3] Gregory Ashton, Moritz Hübner, Paul D. Lasky, Colm Talbot, Kendall Ackley, Sylvia Biscoveanu, Qi Chu, Atul Divakarla, Paul J. Easter, and Boris Goncharov. Bilby: A user-friendly bayesian inference library for gravitational-wave astronomy. *The Astrophysical Journal Supplement Series*, 241(2):27, 2019. Published 2019 April 1. © 2019. The American Astronomical Society. All rights reserved.

[4] Chris Pankow et al. Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817. *Phys. Rev. D*, 98(8):084016, 2018.

[5] R. Abbott et al. GWTC-2: Compact binary coalescences observed by ligo and virgo during the first half of the third observing run. *Phys. Rev. X*, 11:021053, 2021.

[6] R. Abbott, Virgo Collaboration others (LIGO Scientific Collaboration, and KAGRA Collaboration). Gwtc-3: Compact binary coalescences observed by ligo and virgo during the second part of the third observing run. *Phys. Rev. X*, 13:041039, 2023.

[7] 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

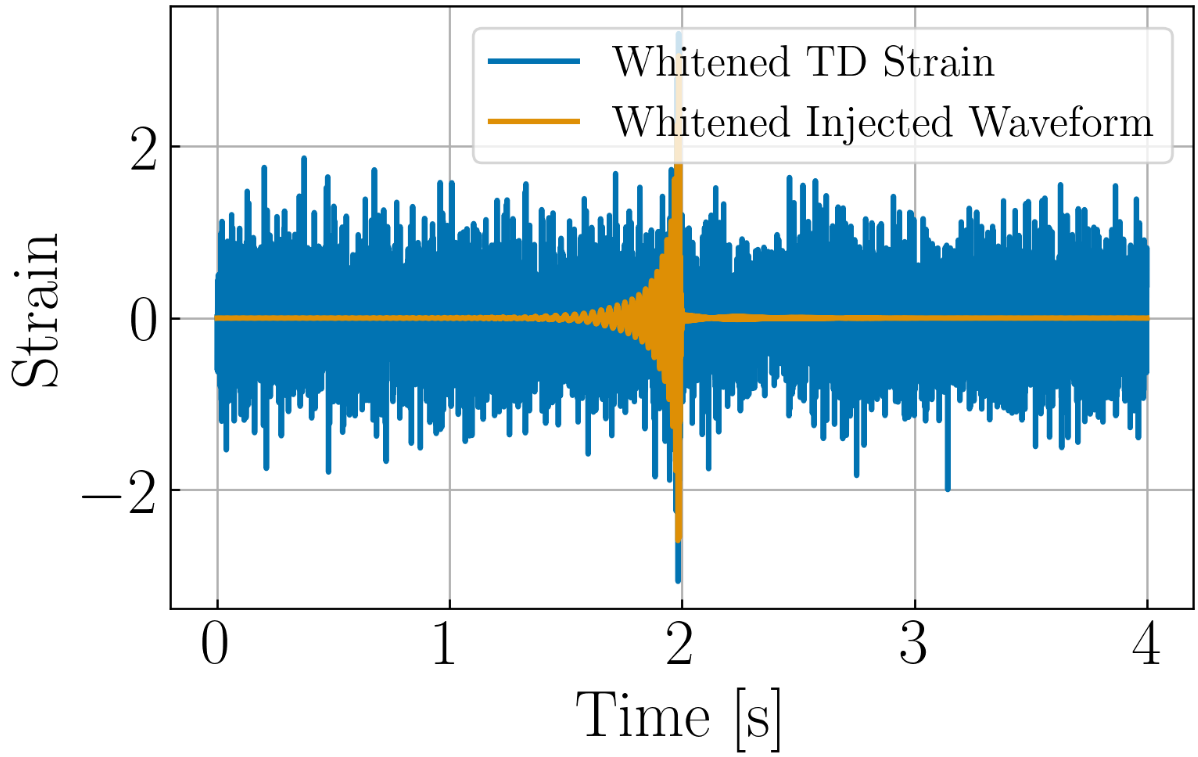

FIG. 2. Figure demonstrating the injected time-domain (TD) strain data. The whitened injected waveform represents the projected waveform onto the detector, while the whitened TD strain data shows the resulting strain observed by the detector.

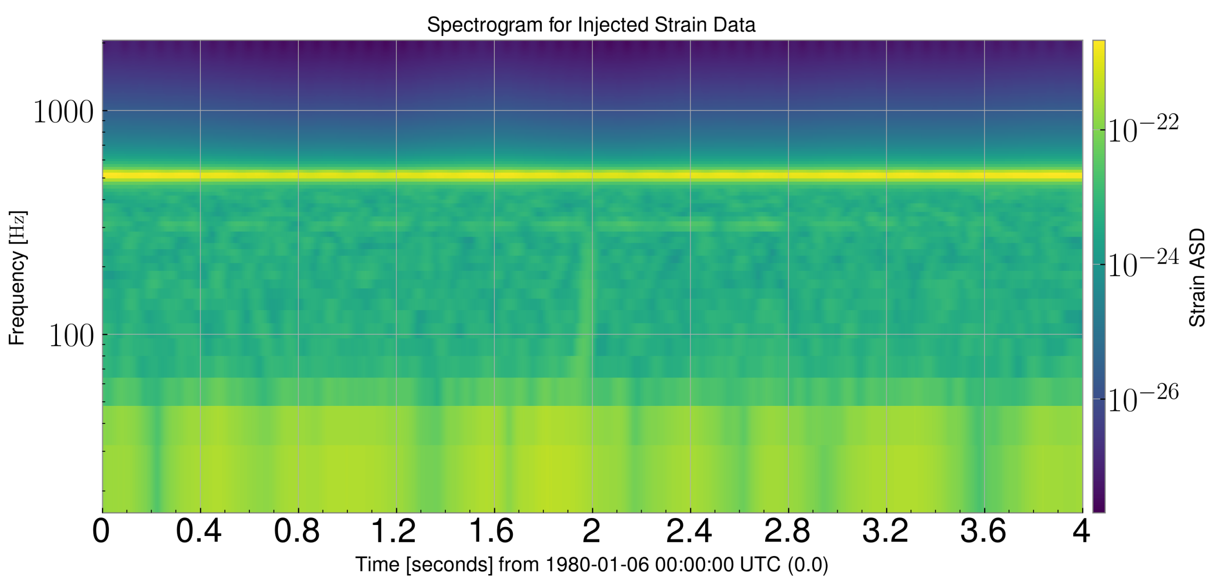

FIG. 3. Figure demonstrating the spectrogram for the injected strain data. The prominent lines in the Amplitude Spectral Density (ASD) mostly align with the noise ASD

the third ligo-virgo observing run. *Classical and Quantum Gravity*, 39(24):245013, 2022.

[8] Neil J. Cornish, Tyson B. Littenberg, Bence Bécsy, Katerina Chatziioannou, James A. Clark, Sudarshan Ghonge, and Margaret Millhouse. Bayeswave analysis pipeline in the era of gravitational wave observations. *Phys. Rev. D*, 103:044006, Feb 2021.

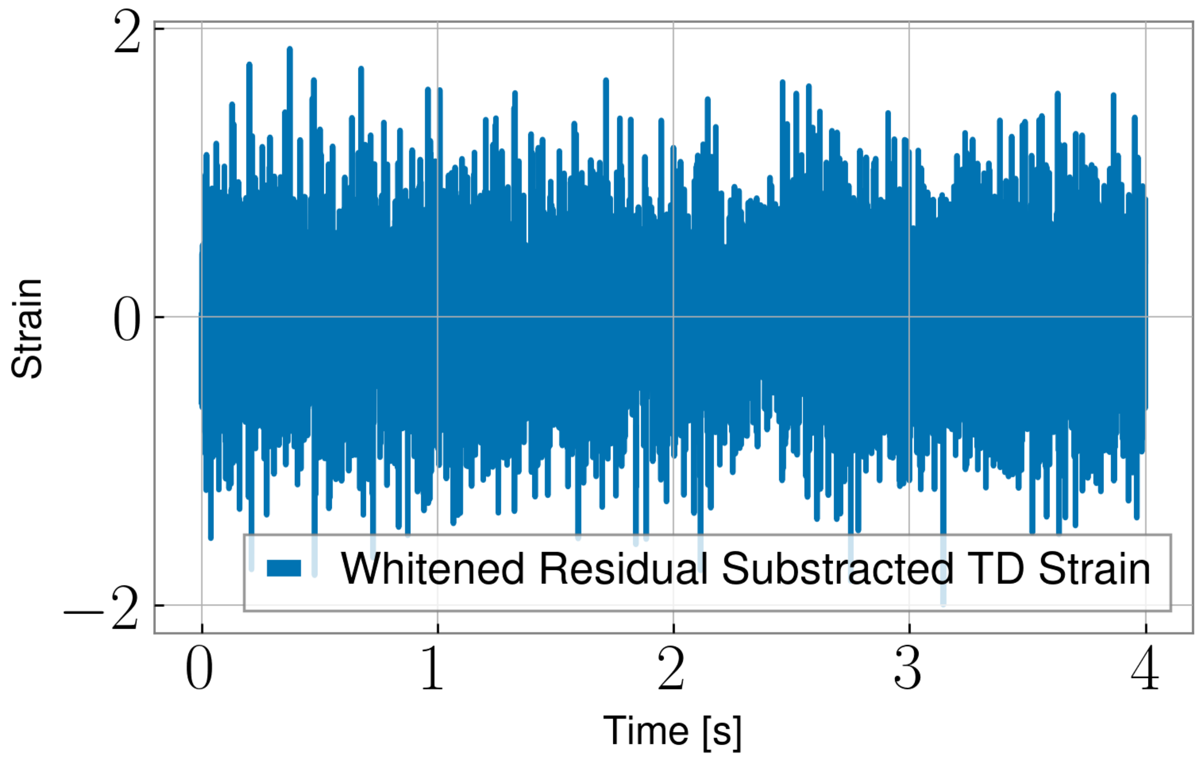

FIG. 4. Figure demonstrating the subtracted TD strain data. The subtraction process involves injecting the negative of the previous waveform. The figure of the residual strain demonstrates agreement with the noise realization of the detector.

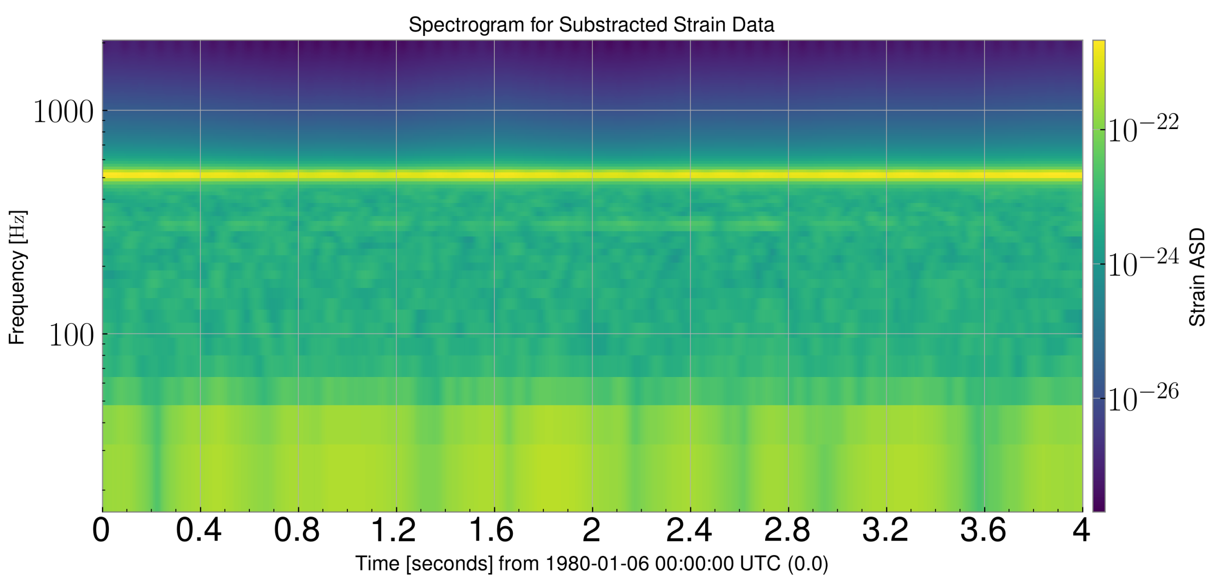

FIG. 5. Figure demonstrating the spectrogram for the subtracted data. The plot reveals subtle differences compared to Figure 3, particularly in the middle section of the spectrogram.

[9] Ethan Payne, Sophie Hourihane, Jacob Golomb, Rhiannon Udall, Derek Davis, and Katerina Chatziioannou. Curious case of GW200129: Interplay between spin-precession inference and data-quality issues. *Phys. Rev. D*, 106:104017, 2022.

[10] Derek Davis, Thomas Massinger, Andrew Lundgren, Jennifer C Driggers, Alex L Urban, and Laura Nuttall. Improving the sensitivity of advanced ligo using noise sub-

traction. *Classical and Quantum Gravity*, 36(5):055011, 2019.

[11] Derek Davis and Marissa Walker. Detector characterization and mitigation of noise in ground-based gravitational-wave interferometers. *Galaxies*, 10(1):12, 2022.

[12] DINGO Collaboration. DINGO: Example toy npe model. https://dingo-gw.readthedocs.io/en/latest/example_toy_npe_model.html, Accessed: 2024-07-15. Accessed on: 2024-07-15.