Sensitive test of non-Gaussianity in gravitational-wave detector data

Ronaldas Macas[*] and Andrew Lundgren[†]
Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, United Kingdom

(Received 16 June 2023; accepted 3 August 2023; published 12 September 2023)

Methods for parameter estimation of gravitational-wave data assume that detector noise is stationary and Gaussian. Real data deviate from these assumptions, which causes bias in the inferred parameters and incorrect estimates of the errors. We develop a sensitive test of non-Gaussianity for real gravitational-wave data that measures meaningful parameters that can be used to characterize these effects. As a test case, we investigate the quality of data cleaning performed by the LIGO-Virgo-KAGRA Collaboration around GW200129, a binary black hole signal that overlapped with the noise produced by the radio frequency modulation. We demonstrate that a significant portion of the non-Gaussian noise is removed below 50 Hz, yet some of the noise still remains after the cleaning; at frequencies above ∼85 Hz, there is no excess noise removed. We also show that this method can quantify the amount of non-Gaussian noise in continuous data, which is useful for general detector noise investigations. To do that, we estimate the difference in non-Gaussian noise in the presence and absence of light scattering noise.

DOI: 10.1103/PhysRevD.108.063016

I. INTRODUCTION

Gravitational-wave (GW) data are often non-Gaussian, meaning it contains outliers. This excess noise is broadly divided into two categories: short and well-defined noise (also called "glitches") and longer-duration broadband excess noise. Both types of excess noise can overlap with real GW signals in time and frequency, therefore affecting GW searches, and source parameter estimation, including the sky localization which is used for electromagnetic follow-up efforts [1–5].

In some cases, the excess noise is recorded by the detector's witness channels. For example, throughout the second LIGO-Virgo observing run, multiple noise couplings were recorded: main power lines and detector calibration lines, as well as the laser beam jitter noise present in LIGO Hanford [6,7].

In principle, recording the noise with witness channels allows us to estimate the noise coupling to the gravitational-wave strain channel and thus remove the noise [8]. With this in mind, the linear noise subtraction method GWSUBTRACT was developed [9]. The method assumes that noise is Gaussian and stationary and that the noise couples linearly. About 30% increase in LIGO Hanford sensitivity was achieved during the second LIGO-Virgo observing run with GWSUBTRACT [9]. There was no noticeable improvement for LIGO Livingston because the "jitter" noise was not present at the detector.

During the third LIGO-Virgo-KAGRA observing run, binary black hole signal GW200129 coincided with the noise caused by the 45 MHz electro-optic modulator driver system at LIGO Livingston [10]. This noise was recorded by multiple radio frequency (rf) channels, one of which was used to clean the data around GW200129 using the previously mentioned linear subtraction method [11]. After the reanalysis of LIGO-Virgo-KAGRA results, Hannam *et al.* found GW200129 to be a highly precessing event; in fact, the measured orbital precession is 10 orders of magnitude higher than any previous weak-field observation [12].

Payne *et al.* questioned whether this event is indeed highly precessing or whether the observed precession is just an artifact of the leftover noise after the data cleaning [13]. While there were various checks performed by the LIGO-Virgo-KAGRA Collaboration to make sure that the linear noise subtraction mitigated the noise [10], there was no statistical estimate of how much rf noise was removed.

In this paper, we present a sensitive statistical test for measuring non-Gaussian noise in GW data. In Sec. II, we derive the normalized Q transform, which allows us to make assumptions about Gaussian data. We also describe how Bayesian statistical modeling can be used to estimate the amount of non-Gaussian noise in GW data. In Sec. III, we apply our method to the data around GW200129 and evaluate the effectiveness of the linear noise subtraction. In addition, we also show that our method can be used as

[*]ronaldas.macas@port.ac.uk
[†]andrew.lundgren@port.ac.uk

an informative real-time noise measurement. Finally, the discussion and conclusions are given in Secs. IV and V, respectively.

II. METHODS

A. Normalized Q transform

GW data containing outliers can be represented by the sum of Gaussian and non-Gaussian time series. In our model, we assume that the Gaussian data follow a Gaussian distribution with a standard deviation of 1, while the non-Gaussian noise is made out of short bursts of transient noise whose amplitudes are drawn from a distribution.

To identify Gaussian and non-Gaussian parts of the time series data, we start with the Q transform defined as

$$X(\tau, f, \mathrm{Q}) = \int_{-\infty}^{+\infty} x(t) w(t-\tau, f, \mathrm{Q}) e^{-i2\pi f t} dt, \quad (1)$$

where $w(t-\tau, f, \mathrm{Q})$ is a window centered on time τ with the width of the window proportional to the quality factor Q,

$$\mathrm{Q} = \frac{f_0}{\Delta f}, \quad (2)$$

where f_0 is central frequency and Δf is the frequency bandwidth [14].

Using Eq. (1), we transform the time series data into time-frequency tiles for certain Q and f_0 values. Contrary to Ref. [15], which uses a variable Q, we chose to use a specific window width for two reasons.

First, we focus only on frequencies where there is non-Gaussianity, which means that our statistical measurement of noise is more precise. Furthermore, by specifying Q and f_0 we normalize the Q tiles in such a way that the average power for Gaussian data in each tile k is 1, i.e.,

$$\frac{1}{\mathrm{N}} \sum_{k=1}^{k=\mathrm{N}} \left[\Re^2(A_k(t)) + \Im^2(A_k(t)) \right] = 1, \quad (3)$$

where $A(t)$ is the amplitude of time series data.

Normalizing the Q transform enables us to make an assumption about the distribution of Q-tile powers for Gaussian data. As shown in Eq. (3), our Q tiles have 2 degrees of freedom: real and imaginary parts of the time-frequency data. As a consequence, we expect that Gaussian data should follow a $\chi^2(2)$ distribution [16], which is simply an exponential distribution with a rate $\lambda = 1/2$ (Fig. 1).

This lets us define a simple statistical test to estimate the Gaussianity of the data: for the number of tiles N, the total Q-tile power for Gaussian data is N and the standard deviation is $\sqrt{4\mathrm{N}}$ (assuming $\mathrm{N} \gg 1$).

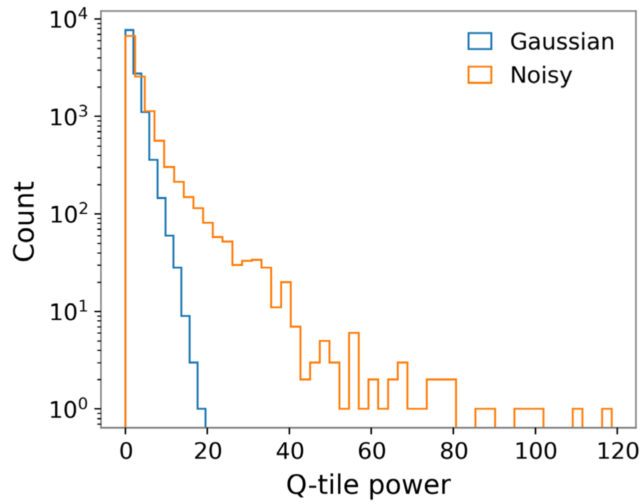

FIG. 1. Q-tile power for $\mathrm{Q} = 8$ and $f_0 = 40$ Hz for 366 s simulated Gaussian data (blue) and 366 s data with excess noise (orange). The distribution of powers for Gaussian data closely follows the expected exponential distribution with a rate $\lambda = 1/2$. Noisy data, however, contain excess noise and therefore do not follow the exponential distribution associated with the Gaussian data.

B. Distribution fitting with Bayesian statistical modeling

For a more informative test of non-Gaussianity, we model the non-Gaussian part of the data. GW data with excess power have two contributions: Gaussian data modeled by $\chi^2(2)$ and a non-Gaussian part represented by an unknown distribution. Data with excess power have much higher energies and a longer tail in contrast with the Gaussian-only data (Fig. 1).

To estimate the amount of non-Gaussianity, we need to fit a mixture of two distributions. We achieve this by using a marginalized mixture model likelihood [17].

The likelihood of a data point z_k belonging to either Gaussian data or non-Gaussian data distribution can be written as

$$p(z_k | q_k, \alpha) = \Pi_{k=1}^{K} [q_k p_1(z_k|\alpha_1) + (1-q_k) p_2(z_k|\alpha_2)], \quad (4)$$

where k is the tile number, q_k is a binary flag, and $p_1(z_k|\alpha_1)$, $p_2(z_k|\alpha_2)$ represents the probability that a point corresponds to the Gaussian noise distribution or a non-Gaussian distribution, respectively.

Prior on q_k is defined as

$$p(q_k) = \begin{cases} F & \text{if } q_k = 0, \\ 1-F & \text{if } q_k = 1, \end{cases} \quad (5)$$

where fraction $F \in [0, 1]$.

Equation (5) allows us to marginalize the likelihood from Eq. (4) giving

$$p(z_k, \alpha) = \Pi_{k=1}^{K}[(1 - F)p_1(z_k|\alpha_1, q_k = 1) + Fp_2(z_k|\alpha_2, q_k = 0)]. \tag{6}$$

To estimate the marginalized likelihood, we use a Bayesian statistical modeling package PyMC with Hamiltonian Monte Carlo "No-U-Turn" sampler (NUTS) [18,19]. Rather than relying on a classical Markov chain Monte Carlo method that does not consider the geometry of parameter space, PyMC with NUTS allows us to recover distribution parameters faster and with fewer samples. It achieves this by specifying a probabilistic model in PyMC, compiling it into a computational graph, providing initial values for parameters, configuring the NUTS sampler, and then running the sampling process. NUTS explores the posterior distribution by proposing and accepting/rejecting parameter values, allowing for efficient and scalable Bayesian inference.

After recovering parameters for each distribution, we estimate the total power in each distribution, i.e., the area under the curve. The ratio of non-Gaussian power to the total power (which we refer to as "fractional power") lets us statistically measure non-Gaussianity in GW detector data.

III. DEMONSTRATION ON REAL DATA

A. Radio frequency noise around GW200129

GW200129 coincided with the 45 MHz radio frequency noise caused by the electro-optic modulator driver system at LIGO Livingston. Luckily, this noise was also recorded by multiple witness channels, thus enabling us to clean the data around GW200129 using the linear subtraction tool GWSUBTRACT [10,20,21].

To estimate the effectiveness of linear subtraction, we compared the data from the original (i.e., not cleaned) LIGO Livingston data frame and the frame that has the excess noise removed using GWSUBTRACT.

First, we identified when rf noise was present in the data around GW200129 using Z score. Out of 4096 s, we selected in total 366 s considered to contain radio frequency noise. After that, the data subset was Q tiled for various Q and f_0 values.

For the first test of non-Gaussianity, we simply estimate what is the average tile power for both the original and the GWSUBTRACT data. To do that, we Q tile the data for various Q and f_0 values, which gives us the total number of tiles and the power in each tile. Since we use the normalized Q tiles, the average tile power for Gaussian data is 1. Figure 2(a) shows the corresponding average tile power for quality factor Q = 8. We chose the Q value of 8 in this and the following figures because low Q-value tiles fit the short-duration glitches such as rf noise better.

As a more advanced test, we want to measure the amount of non-Gaussianity. In order to do that, we model the data with an exponential distribution with a rate $\lambda = 1/2$ and a half Student T distribution [22] with probability density function given by

$$f(x|\sigma, \nu) = \frac{2\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\nu\pi\sigma^2}}\left(1 + \frac{1}{\nu}\frac{x^2}{\sigma^2}\right)^{-\frac{\nu+1}{2}}, \tag{7}$$

where ν is degrees of freedom, σ is the normality parameter, and Γ is the gamma function. We chose half Student T distribution because it has a long tail that closely matches the non-Gaussian data (see, for example, Fig. 1).

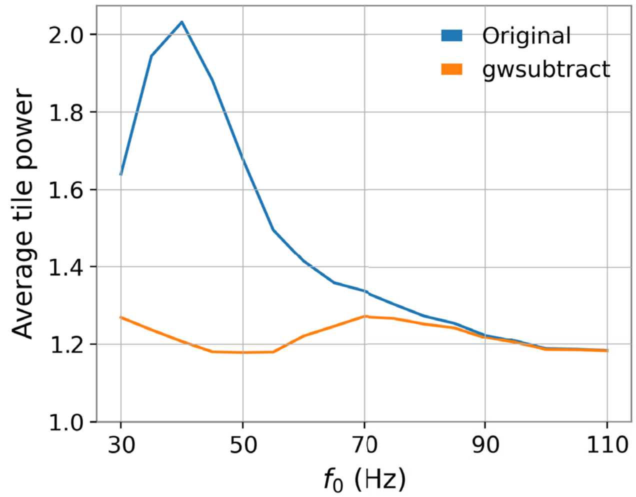

(a) Average tile power for data around GW200129 with quality factor $Q = 8$.

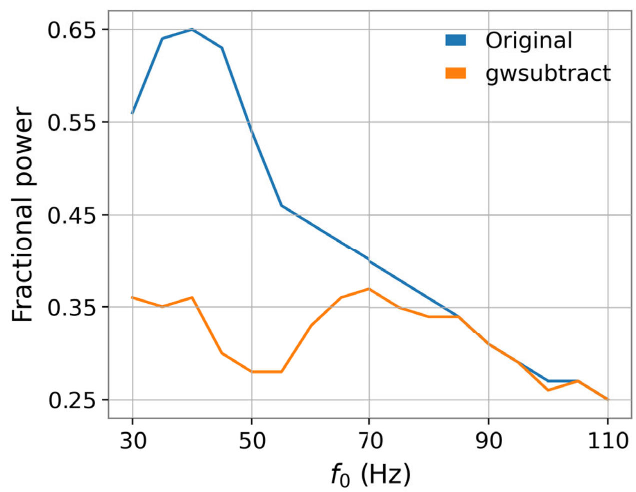

(b) Fractional power for data around GW200129 with quality factor $Q = 8$.

FIG. 2. Average tile power (a) and fractional power (b) for data around GW200129 with quality factor Q = 8. Linear subtraction GWSUBTRACT removes a significant portion of the non-Gaussian noise below 50 Hz, yet some of the noise still remains after the cleaning. At frequencies above ~85 Hz, there is no excess noise removed.

After recovering distribution parameters, we estimate the fractional power, which is the ratio of non-Gaussian power to the total power. Figure 2(b) shows the fractional power for GW200129.

B. Light scattering noise

We also apply both tests to measure the excess noise in the time series data that is continuous. For this example, we selected 10 min of data with a relatively low amount of light scattering noise and 10 min of scattering-heavy data recorded by LIGO Livingston on January 6, 2020. Light scattering glitches are caused by stray light reflection in the beam tube, which can happen due to excessive ground motion [23–25]. As a result, such noise transients can last up to minutes with a typical frequency range of 20–50 Hz.

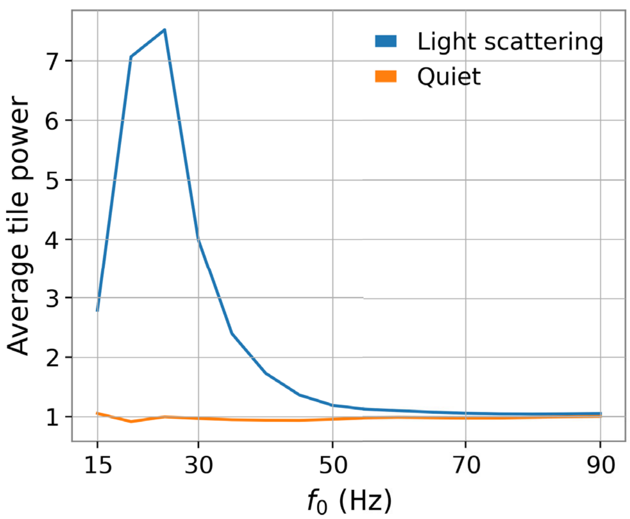

(a) Average tile power for data with quality factor $Q = 8$ in the presence and absence of light scattering artifacts.

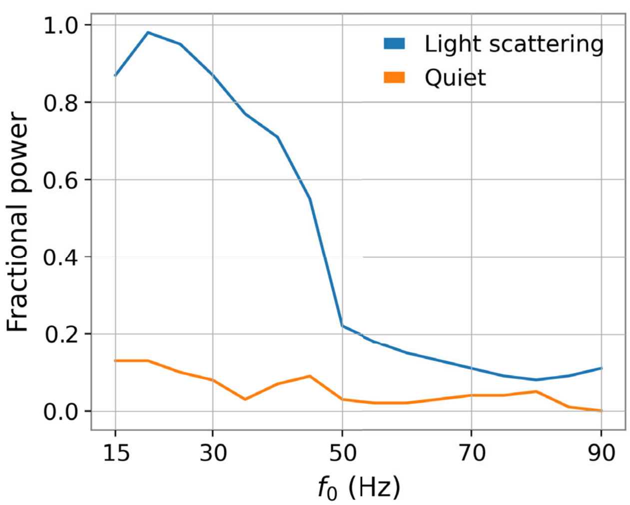

(b) Fractional power for data with quality factor $Q = 8$ in the presence and absence of light scattering artifacts.

FIG. 3. Average tile power (a) and fractional power (b) for data with quality factor Q = 8. Time series data in blue have excess power at the 15–50 Hz range caused by the light scattering artifacts, whereas the time series data in orange with no light scattering artifacts have a relatively small amount of non-Gaussian noise across all frequencies.

We use the same procedure as in the previous example of GW200129. First, we calculate the Q-tile power for specific Q and f_0 values giving us the average tile power, which is equal to 1 for Gaussian data. Then, we fit the Q-tile power with $\chi^2(2)$ and half Student T distributions to estimate the fractional power, which is the ratio of non-Gaussian power to the total power. Figure 3 shows the average tile power and the fractional power for both low noise and scattering-heavy time series using quality factor Q = 8.

IV. DISCUSSION

In this paper, we present a method to estimate the amount of non-Gaussian noise in the data based on the normalized Q transform. This method has multiple advantages over other statistical tests. While comparing power spectral density (PSD) between the cleaned and original data provides a quantitative measurement of the removed power, it does not discriminate whether the removed power is non-Gaussian, i.e., the excess noise, or Gaussian [7]. Furthermore, the PSD estimation is usually performed over a long duration, e.g., 512 s, meaning that such a test becomes less sensitive to removing short-duration noise bursts like in the data around GW200129.

Contrary to the PSD comparison, our test discriminates between the Gaussian and non-Gaussian noise; we also analyze only the times when the rf noise is present, which makes our test much more sensitive than the PSD comparison.

Another commonly used method to estimate how well the data are cleaned relies on the visual inspection of spectrograms [7]. However, this becomes problematic when the excess noise is weak and is comparable to Gaussian noise. This issue becomes even more precarious when the excess noise overlaps a GW signal, as in the case of GW200129. Conversely, our method provides a quantitative measurement of the excess noise.

By using normalized Q tiles, we put a constraint that Gaussian data have an average tile power of 1, whereas the data with excess noise are expected to have more power on average. Furthermore, we found that the Gaussian data must follow a $\chi^2(2)$ distribution after we normalize the Q tiles. As a result, the data containing Gaussian and non-Gaussian noise can be decomposed into two distributions, one of which is $\chi^2(2)$.

To estimate the amount of non-Gaussian noise in GW data, we fit the normalized Q-transform power using Bayesian statistical modeling with the Hamiltonian Monte Carlo algorithm. By doing this, we reconstruct the parameters of both Gaussian and non-Gaussian distributions.

As the first application of this method, we investigate if (and how much) GWSUBTRACT cleaned the data around GW200129. As shown in Fig. 2, rf noise is mostly affecting the data in the 30–50 Hz range, where the non-Gaussian power is at least 55% of the total data power, i.e., the frequencies where precession effects are especially important for GW200129 [13]. Linear subtraction GWSUBTRACT removes up to 30% of the fractional power in this frequency range, indicating that it does remove the rf noise, albeit not entirely.

At higher frequencies, there is less rf noise in the data, which could explain why GWSUBTRACT is less efficient at these frequencies compared to the 30–50 Hz range. Above $\sim$85 Hz, GWSUBTRACT does not remove any noise, which is reasonable given that the witness channel used to clean the data does not have any information above this frequency.

Our simpler test showing the average tile power agrees well with the fractional power results (Fig. 2). Both frames, original and GWSUBTRACT, have more power per tile than the Gaussian data. However, the original data have considerably more power in the 30–50 Hz range (up to 70% more).

We also tested if our method can be applied as a live monitoring tool of excess noise in GW detector data. For that, we compared 10 min of relatively quiet data with 10 min of data containing many scattering glitches at LIGO Livingston on January 6, 2020.

We found that the fractional power is much higher for data with many scattering glitches than for the quiet data [Fig. 3(b)]. The excess noise is concentrated within 10–50 Hz and peaks at 20 Hz, where the non-Gaussian noise power is 0.96 of the total power. The frequency of light scattering noise (20–50 Hz) is similar to the excess noise at 15–50 Hz, indicating that our test correctly identified the presence of excess noise. For 10 min of relatively low noise, the fractional power is much lower and stays relatively constant across all frequencies.

Similarly, our average tile power test indicates that the data with scattering glitches contain up to $\sim$7 times more power than the Gaussian data [Fig. 3(a), blue]. Data with no scattering glitches [Fig. 3(a), orange] stay relatively constant across all frequencies and have an average tile power close to Gaussian data, i.e., 1.

V. CONCLUSIONS

Gravitational-wave data contain many noise transients that can affect GW searches and estimation of source parameters such as the precession of a binary black hole. It is possible to mitigate excess noise from GW data in some cases, for example, by correlating information from noise witness channels. However, the lack of meaningful and sensitive statistical tests prevents estimating the effectiveness of such noise mitigation tools.

In this paper, we propose a novel method to statistically measure the non-Gaussian noise in GW detector data. By using normalized Q transform, we constrain how Gaussian data should be distributed, which allows us to determine the amount of non-Gaussian noise in the data.

Our method provides two statistical tests to measure the non-Gaussianity. The first one measures the average tile power which, for Gaussian data, must be equal to 1. The second test estimates the relative power of non-Gaussian data compared to the total power, which we define as fractional power. Both our tests are sensitive to excess noise changes in frequency, which allows for a more precise noise measurement than averaging across all frequencies.

In order to show the effectiveness of our statistical tests, we explore two scenarios. First, we investigate how well GWSUBTRACT, a linear subtraction tool, removes the noise around the binary black hole event GW200129. We find that GWSUBTRACT removes a significant portion of the non-Gaussian noise at lower frequencies (30–70 Hz), yet it fails to remove the non-Gaussian noise completely. There is also no noise removed above $\sim$85 Hz.

In the second scenario, we test if our method can be applied as a live monitoring tool of excess noise in GW detector data. We show that for data that have many light scattering noise artifacts, our method correctly identifies the excess noise within the light scattering noise frequency range (20–50 Hz).

ACKNOWLEDGMENTS

We thank Marissa Walker and members of the LIGO Detector Characterization group for their valuable input during the preparation of this manuscript. We also thank the anonymous referee for valuable comments which improved the manuscript during the review. R. M. is supported by STFC Grants No. ST/S000550/1, No. ST/T000325/1, No. ST/V005715/1, and No. ST/X002225/1. The authors are grateful for computational resources provided by the LIGO Laboratory and supported by National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459. This research has made use of data, software and/or web tools obtained from the Gravitational Wave Open Science Center [26], a service of LIGO Laboratory, the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This material is based upon work supported by NSF's LIGO Laboratory which is a major facility fully funded by the National Science Foundation. LIGO Laboratory and Advanced LIGO are funded by the U.S. National Science Foundation (NSF) as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. Additional support for Advanced LIGO was provided by the Australian Research Council. Virgo is funded, through the European Gravitational Observatory (EGO), by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale di Fisica Nucleare (INFN), and the Dutch Nikhef, with contributions by institutions from Belgium, Germany, Greece, Hungary, Ireland, Japan, Monaco, Poland, Portugal, Spain. KAGRA is supported by Ministry of Education, Culture, Sports, Science and Technology (MEXT), Japan Society for the Promotion of Science (JSPS) in Japan; National Research Foundation (NRF) and Ministry of Science and ICT (MSIT) in Korea; Academia Sinica (AS) and National Science and Technology Council (NSTC) in Taiwan. Various parts of the analysis used GWpy [27], PyCBC [28,29] PyMC [18], NumPy [30], SciPy [31], IPython [32], JUPYTER NOTEBOOK [33], PANDAS [34], MATPLOTLIB [35], ARVIZ [36], and CORNER [37].

[1] T. D. Canton, S. Bhagwat, S. V. Dhurandhar, and A. Lundgren, Effect of sine-Gaussian glitches on searches for binary coalescence, Classical Quantum Gravity **31**, 015016 (2013).

[2] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO's first observing run, Classical Quantum Gravity **35**, 065010 (2018).

[3] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, Classical Quantum Gravity **37**, 055002 (2020).

[4] C. Pankow, K. Chatziioannou, E. A. Chase, T. B. Littenberg, M. Evans, J. McIver, N. J. Cornish, C.-J. Haster, J. Kanner, V. Raymond *et al.*, Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817, Phys. Rev. D **98**, 084016 (2018).

[5] R. Macas, J. Pooley, L. K. Nuttall, D. Davis, M. J. Dyer, Y. Lecoeuche, J. D. Lyman, J. McIver, and K. Rink, Impact of noise transients on low latency gravitational-wave event localization, Phys. Rev. D **105**, 103021 (2022).

[6] R. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), GWTC-1: A Gravitational-Wave Transient Catalog of Compact Binary Mergers Observed by LIGO and Virgo During the First and Second Observing Runs, Phys. Rev. X **9**, 031040 (2019).

[7] Derek Davis *et al.*, LIGO detector characterization in the second and third observing runs, Classical Quantum Gravity **38**, 135014 (2021).

[8] Nguyen *et al.*, Environmental noise in Advanced LIGO detectors, Classical Quantum Gravity **38**, 145001 (2021).

[9] D. Davis, T. Massinger, A. Lundgren, J. C. Driggers, A. L. Urban, and L. Nuttall, Improving the sensitivity of Advanced LIGO using noise subtraction, Classical Quantum Gravity **36**, 055011 (2019).

[10] R. Abbott *et al.* (The LIGO Scientific Collaboration and the Virgo Collaboration and the KAGRA Collaboration), GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo During the Second Part of the Third Observing Run, arXiv:2111.03606 [Phys. Rev. X (to be published)].

[11] D. Davis, T. B. Littenberg, I. M. Romero-Shaw, M. Millhouse, J. McIver, F. D. Renzo, and G. Ashton, Subtracting glitches from gravitational-wave detector data during the third LIGO-Virgo observing run, Classical Quantum Gravity **39**, 245013 (2022).

[12] M. Hannam *et al.*, General-relativistic precession in a black-hole binary, Nature (London) **610**, 652 (2022).

[13] E. Payne, S. Hourihane, J. Golomb, R. Udall, D. Davis, and K. Chatziioannou, Curious case of GW200129: Interplay between spin-precession inference and data-quality issues, Phys. Rev. D **106**, 104017 (2022).

[14] S. K. Chatterji, The search for gravitational wave bursts in data from the second LIGO science run (2005), https://dspace.mit.edu/handle/1721.1/34388.

[15] L. Vazsonyi and D. Davis, Identifying glitches near gravitational-wave signals from compact binary coalescences using the q-transform, Classical Quantum Gravity **40**, 035008 (2023).

[16] W. G. Cochran, The χ^2 test of goodness of fit, Ann. Math. Stat. **23**, 315 (1952).

[17] D. Foreman-Mackey, Blog post: Mixture models, 10.5281/zenodo.15856 (2014).

[18] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, Probabilistic programming in PYTHON using PyMC3, PeerJ Comput. Sci. **2**, e55 (2016).

[19] M. D. Hoffman and A. Gelman, The no-u-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, arXiv:1111.4246.

[20] D. Davis *et al.*, Data quality vetoes applied to the analysis of LIGO data from the third observing run (2021), https://dcc.ligo.org/LIGO-T2100045/public.

[21] A. Lundgren, aLIGO LLO Logbook, 47707 (2019), https://alog.ligo-la.caltech.edu/aLOG/index.php?callRep=47707.

[22] B. C. Sutradhar, On the characteristic function of multivariate student t-distribution, Can. J. Stat./La Revue Canadienne de Statistique **14**, 329 (1986).

[23] Soni *et al.*, Reducing scattered light in LIGO's third observing run, Classical Quantum Gravity **38**, 025016 (2020).

[24] Accadia *et al.*, Noise from scattered light in Virgo's second science run data, Classical Quantum Gravity **27**, 194011 (2010).

[25] A. E. Tolley, G. S. Cabourn Davies, I. W. Harry, and A. P. Lundgren, ArchEnemy: Removing scattered-light glitches from gravitational wave data, Classical Quantum Gravity **40**, 165005 (2023).

[26] https://www.gw-openscience.org.

[27] D. M. Macleod, J. S. Areeda, S. B. Coughlin, T. J. Massinger, and A. L. Urban, GWpy: A Python package for gravitational-wave astrophysics, SoftwareX **13**, 100657 (2021).

[28] G. S. Davies, T. Dent, M. Tápai, I. Harry, C. McIsaac, and A. H. Nitz, Extending the PyCBC search for gravitational

waves from compact binary mergers to a global network, Phys. Rev. D **102**, 022004 (2020).

[29] T. D. Canton, A. H. Nitz, B. Gadre, G. S. Davies, V. Villa-Ortega, T. Dent, I. Harry, and L. Xiao, Realtime search for compact binary mergers in Advanced LIGO and Virgo's third observing run using PyCBC live, Astrophys. J. **923**, 254 (2021).

[30] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, Array programming with NumPy, Nature (London) **585**, 357 (2020).

[31] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in PYTHON, Nat. Methods **17**, 261 (2020).

[32] F. Perez and B. E. Granger, IPython: A system for interactive scientific computing, Comput. Sci. Eng. **9**, 21 (2007).

[33] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, Jupyter notebooks—a publishing format for reproducible computational workflows, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, edited by F. Loizides and B. Schmidt (IOS Press, Amsterdam, 2016), pp. 87–90.

[34] W. McKinney *et al.*, Data structures for statistical computing in PYTHON, in *Proceedings of the 9th PYTHON in Science Conference* (Austin, TX, 2010), Vol. 445, pp. 51–56.

[35] J. D. Hunter, MATPLOTLIB: A 2d graphics environment, Comput. Sci. Eng. **9**, 90 (2007).

[36] R. Kumar, C. Carroll, A. Hartikainen, and O. Martin, Arviz a unified library for exploratory analysis of Bayesian models in PYTHON, J. Open Source Software **4**, 1143 (2019).

[37] D. Foreman-Mackey, CORNER.PY: Scatterplot matrices in PYTHON, J. Open Source Software **1**, 24 (2016).