

Detecting and diagnosing terrestrial gravitational-wave mimics through feature learning

Robert E. Colgan^{1,2}, Zsuzsa Márka⁵, Jingkai Yan,^{2,3} Imre Bartos,⁴ John N. Wright,^{2,3} and Szabolcs Márka⁶

¹*Department of Computer Science, Columbia University in the City of New York,
500 W. 120th Street, New York, New York 10027, USA*

²*Data Science Institute, Columbia University in the City of New York,
550 W. 120th Street, New York, New York 10027, USA*

³*Department of Electrical Engineering, Columbia University in the City of New York,
500 W. 120th Street, New York, New York 10027, USA*

⁴*Department of Physics, University of Florida, PO Box 118440, Gainesville, Florida 32611-8440, USA*

⁵*Columbia Astrophysics Laboratory, Columbia University in the City of New York,
538 W. 120th Street, New York, New York 10027, USA*

⁶*Department of Physics, Columbia University in the City of New York,
538 W. 120th Street, New York, New York 10027, USA*



(Received 6 July 2022; accepted 15 February 2023; published 20 March 2023)

As engineered systems grow in complexity, there is an increasing need for automatic methods that can detect, diagnose, and even correct transient anomalies that inevitably arise and can be difficult or impossible to diagnose and fix manually. Among the most sensitive and complex systems of our civilization are the detectors that search for incredibly small variations in distance caused by gravitational waves—phenomena originally predicted by Albert Einstein to emerge and propagate through the universe as the result of collisions between black holes and other massive objects in deep space. The extreme complexity and precision of such detectors causes them to be subject to transient noise issues that can significantly limit their sensitivity and effectiveness. They are also subject to nearly constant development, improvement, commissioning and other invasive actions that change the nature of the data and its artifact contamination. In this work, we present a demonstration of a method that can detect and characterize emergent transient anomalies of such massively complex systems. We illustrate the performance, precision, and adaptability of the automated solution via one of the prevalent issues limiting gravitational-wave discoveries: noise artifacts of terrestrial origin that contaminate gravitational wave observatories’ highly sensitive measurements and can obscure or even mimic the faint astrophysical signals for which they are listening. Specifically, we demonstrate how a highly interpretable convolutional classifier can automatically learn to detect transient anomalies from auxiliary detector data without needing to observe the anomalies themselves. We also illustrate several other useful features of the model, including how it performs automatic variable selection to reduce tens of thousands of auxiliary data channels to only a few relevant ones; how it identifies behavioral signatures predictive of anomalies in those channels; and how it can be used to investigate individual anomalies and the channels associated with them. The solution outlined is broadly applicable, enabling automated anomaly discovery and characterization and human-in-the-loop anomaly elimination.

DOI: [10.1103/PhysRevD.107.062006](https://doi.org/10.1103/PhysRevD.107.062006)

I. INTRODUCTION

Gravitational-wave detectors such as KAGRA, Virgo, GEO600, and Laser Interferometer Gravitational-wave Observatory (LIGO) use highly sensitive interferometers [1–4] to measure the extraordinarily tiny variations gravitational waves cause as they pass through the detectors. Discoveries [5–7] of gravitational-wave detectors are documented in extensive catalogs, such as GWTC-1 [8], GWTC-2 [9], GWTC-2.1 [10], GWTC-3 [11], 1-OGC [12], 2-OGC [13], 3-OGC [14], IAS-Princeton [15–17],

and eccentric localizations [18]. Due to the detectors’ precision, they are also subject to various forms of noise that can appear in the gravitational-wave measurement data and interfere with their ability to detect and characterize gravitational waves. Among the most troublesome forms of noise are short, loud anomalies known within LIGO as “glitches,” which can obscure or even mimic real gravitational waves. Unlike ever-present background noise [19,20], glitches cannot simply be subtracted from the desired signal. They have a wide variety of causes, some of

which are well understood and many of which are not. Attempts to mitigate various types of glitches have been a significant focus of LIGO’s engineering efforts for more than a decade, and have included multiple machine learning–based approaches [20–48]. In addition to the main interferometer data channel sensitive to gravitational waves, known as the “strain,” each LIGO detector continuously records hundreds of thousands of channels of “auxiliary” data describing various aspects of the detector’s state and internal and external environment. Although the vast majority of these channels do not record useful information for identifying the source of specific glitches, some have been found to measure behavior predictive of or associated with certain types of glitches. Understanding which channels may be associated with glitches and how is therefore an important avenue for diagnosing the root causes of glitches and in some cases reducing or eliminating them.

Recent machine learning–based approaches to glitch understanding and mitigation include works that investigate the association between glitches and LIGO’s auxiliary data channels by casting glitch detection as a classification problem using only nonastrophysically sensitive auxiliary channel data as input [21,22,27]. Generally, glitch detection methods, including those referenced above, may identify times likely to contain glitches using auxiliary channels that have previously been observed to correlate with noise transients in the strain data, and also often directly analyze the gravitational-wave strain, sometimes in combination with auxiliary data [24–26,29,33,34,36–47,49–54]. However, machine learning–based methods that use strain data only indirectly (as a labeled dataset) during training to learn correlations between glitches and auxiliary channels can provide a model that produces a strain-independent corroboration for glitch probability [21,22,27]. Because such models base their predictions only on astrophysically insensitive auxiliary data, they present a lower risk of inadvertently labeling an astrophysical event as a glitch. Notably, achieving a sufficient level of accuracy with a strictly auxiliary channel–based model necessarily implies that it has recognized and learned to be sensitive to behavior observed in only the auxiliary channels that is predictive of glitches in the gravitational-wave strain data stream. These automatically identified associations can then be analyzed by detector domain experts for potentially unexplored insights into the origins and causes of glitches. This pursuit could ultimately lead to mitigations that reduce glitch frequency, increase the sensitivity of the detector, and enable sufficiently confident detections of astrophysical events that would otherwise have been rejected as too uncertain [22].

In this work, we demonstrate and apply one such model to a real dataset of a certain type of glitch that was particularly prevalent at LIGO Livingston Observatory (LLO) at various times during LIGO’s Observing Run

3b (O3b). We demonstrate the model’s ability to identify such glitches with a remarkable accuracy greater than 97%. We also illustrate clearly how it arrives at its predictions and highlight several other useful aspects of the model, including its high degree of interpretability and the insights it can provide into the glitches and channels analyzed. The results provide a demonstration of the capabilities of the models described. We hope they will stimulate further exploration and application to additional glitch types and datasets, eventually leading to breakthroughs that allow glitches to be reduced or eliminated and the detectors’ sensitivity to be increased—which could enable sufficiently confident detections of more gravitational waves and further our understanding of our universe.

II. FEATURE LEARNING FOR TRANSIENT ANOMALIES

Machine learning models have been demonstrated in recent works to be able to learn to predict the presence or absence of glitches in high-dimensional gravitational-wave astronomy data by considering only auxiliary channel information, without looking at the gravitational-wave strain data stream in which the glitches themselves appear [21,22,27]. In this section, we describe one such model initially proposed in [21], which we refer to as LF. The model is highly interpretable and well-suited to efficient strain-independent glitch detection, offering several useful features including:

- (i) Strain-independent glitch prediction. It is extremely important for potential gravitational-wave detections to be rigorously vetted to ensure they are in fact true astrophysical events. Terrestrial-origin glitches can obscure or mimic gravitational waves, so it is also important to identify and distinguish them from astrophysical signals with high confidence. Unlike existing glitch-detection methods such as Omicron [55] and GravitySpy [34] that directly analyze the gravitational-wave strain data stream to search for high-power transient noise events, the methods discussed here consider only the detector’s auxiliary data channels. As such, they provide independent corroboration from an entirely different data source; for example, because auxiliary channels are generally not astrophysically sensitive (and we take pains to remove any that might be from consideration—see Sec. III B), it is unlikely that these methods would inadvertently identify a true gravitational wave as a glitch. [Sec. IV A].
- (ii) Channel selection. The models are trained on tens of thousands of auxiliary channels. Certain channels are known to detector experts to monitor aspects of the detector’s behavior associated with glitches, but it would be impossible to comprehensively analyze all channels manually. Absent such domain expert knowledge and time, machine learning models can

discover connections between auxiliary channels and glitches that might otherwise have gone undetected. Notably, the models discussed here are explicitly tuned to learn to ignore input channels that are not useful for making its predictions—the vast majority of them, in this setting. Such a practice has previously been demonstrated (e.g., in [21,22]) to improve both interpretability and accuracy. These methods achieve excellent performance while ignoring all but a few dozen to few hundred channels, which can then be further investigated manually by detector domain experts. [Sec. IV B].

- (iii) Learned features. In addition to simply telling us which channels to look at, the models discussed here also learn behavioral signatures that correspond to patterns of behavior associated with glitches (or the lack thereof) in individual channels. For example, if a certain type of glitch is frequently preceded by a level change in a particular channel, we can expect the feature corresponding to that channel to learn to be sensitive to such a pattern. A few of these learned features are illustrated in Figs. 1 and 11. [Sec. IV C].
- (iv) Channel contributions to individual glitches. For any individual glitch (or set of glitches) of interest, we can feed the model the auxiliary channel data for those glitches and examine the contributions of each channel to the classifier’s result for those glitches. This could be useful if, for example, there are multiple potential causes of a given glitch and we need to determine which one is most likely at play. [Sec. IV D].

Other algorithms including Omicron [55], iDQ [27], and hveto [43] also analyze auxiliary channels for glitch source identification; thus the qualities specified above may not be necessary unique to the LF model of [21].

For the experiments presented here, we employ a slightly modified version of the LF model introduced in [21]. The model is a flat convolutional binary classifier parametrized by a learned filter $f_p \in \mathbb{R}^T$ of desired length T^1 for each of the P input time series (auxiliary channels) as well as a single learned bias term b_{s_p} for all channels of each sample rate s_p .²

¹We set the filter length T to 96 samples for all channels based on the finding of [21] that six seconds of data for 16-Hz channels performed well. As discussed in Sec. A, for computational efficiency we used the same number of samples for channels regardless of sample rate, so the model sees six seconds of data for a 16-Hz channel but only 0.375 seconds of data for a 256-Hz channel, for example.

²The minor modification we make to the LF model of [21] is the introduction of a separate bias term b_{s_p} for each sample rate s_p rather than a single bias term for all channels, as we found this slightly improved performance. ([21] used only 16-Hz channels in its experiments.)

To make a prediction at time t , the model takes as input the normalized data (centered at t) for all channels $x_p[t] \in \mathbb{R}^T$ of each sample rate s_p ; cross-correlates the data with the corresponding filter f_p ; sums the results for each sample rate s_p and adds the corresponding bias term b_{s_p} ; and sums the results across all sample rates. The result is then translated into a probability estimate between 0 and 1 (where 0 indicates the model is certain no glitch is present and 1 indicates it is certain there is a glitch) by passing it through a sigmoid function. Mathematically,

$$\hat{y}_t = \sigma \left(\sum_{s_p} \left(\sum_p f_p \star x_p[t] \right) + b_{s_p} \right), \quad (1)$$

where σ denotes the logistic function $\sigma(x) = (1 + \exp(-x))^{-1}$ and \star denotes discrete correlation.

As in [21], we employ a sparsifying regularizer—the elastic net [56]—during training to encourage most of the filters to be 0, i.e., $\|f_p\| = 0$ for most p , because sparsity was found to be essential to good performance as well as beneficial to interpretability [21,22].

After training, the remaining nonzero filters correspond to the channels the model finds useful in making its predictions, indicating channels with behavior that was associated with the appearance of glitches or lack thereof during the training period. The magnitude of the filters also provides an indication of the relative importance of each channel to the model’s predictions (intuitively, when correlated with a higher-magnitude filter, an input data segment will contribute more heavily to the sum and resulting probability estimate than the same segment correlated with a lower-magnitude filter) [21].

In Figs. 1 and 2, we illustrate visually how the model arrives at its predictions given a potential glitch time t . Figure 1 shows two examples of learned filters f_p . Figure 2 illustrates two channels’ data $x_p[t]$ (blue) during several glitches around GPS time $t = 1, 259, 253, 345$ (the center of the x-axis) as well as the cross-correlation (orange) between the channel data centered at each point on the x-axis and the corresponding filter illustrated in Fig. 1. To arrive at an estimate that there is a glitch present at that time, the model would sum the cross-correlations at that time over all channels considered. In Fig. 2, the cross-correlation values for both channels at that time are positive, so both channels are contributing to the model predicting the presence of a glitch. The left channel’s cross-correlation is significantly greater, however, so it makes a larger contribution to the model’s confidence that a glitch is present.

III. APPLICATION TO SCATTERED LIGHT GLITCHES

To illustrate how this method can be employed to interrogate a particular variety of recurring glitch, in this

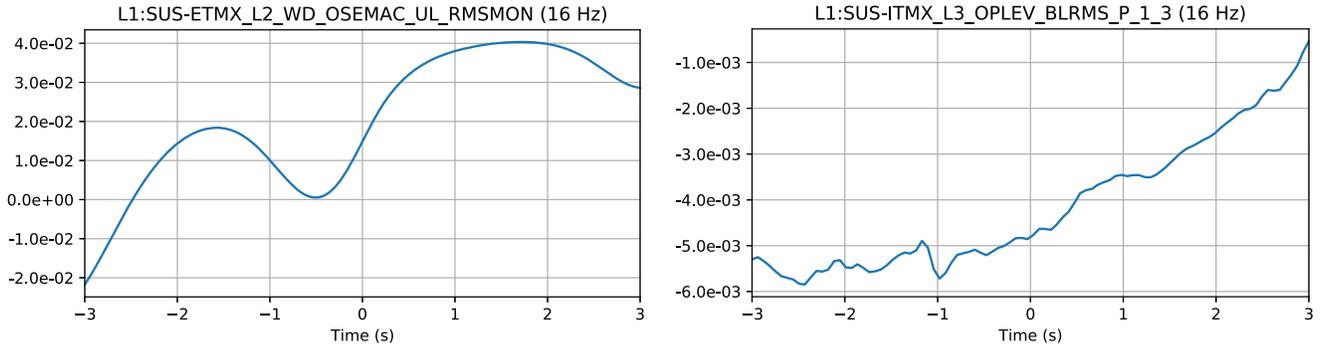


FIG. 1. An example of two learned filters in a trained LF model corresponding to two of the channels found by the model to contain behavior associated with scattered light glitches. See Sec. IV for more details about what the channels in question measure and what behaviors the model has learned to be sensitive to.

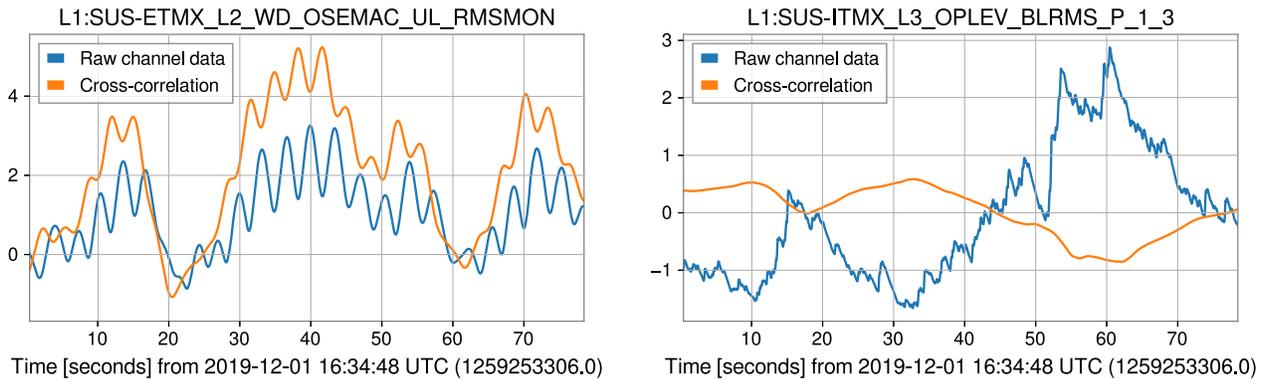


FIG. 2. Blue: normalized channel data from the two channels whose corresponding filters are illustrated in Fig. 1 over time during several scattered light glitches, including a 19.75-second-long one centered at GPS time 1,259,253,341.438 (approximately 35.4 seconds on the x-axis), a 9.5-second-long one centered at 1,259,253,318.188 (12.2 seconds), and several shorter ones from 1,259,253,370.438 to 1,259,253,382.719 (64.4 to 76.7 seconds). The y-axis corresponds to standard deviations above or below the mean (computed over all samples during the training period). Orange: the result of the cross-correlation between the filters shown in Fig. 1 and the raw channel data over time. At each point along the x-axis, the filter is cross-correlated with the corresponding number of samples of raw data centered at that point; the result is plotted on the y-axis. Note that the filter for the channel on the right is negative, so when the raw value drops below 0 the cross-correlation is positive. Figure 7 shows a visualization of the strain data during the same time period illustrated above—in which the glitches are clearly visible—as well as the model’s final glitch probability estimate over that period.

section we demonstrate its application to a particularly troublesome glitch type known as “scattered light” [57–64] during a period when the LLO detector was experiencing an elevated frequency of this type of glitch. We chose this particular class of glitch to validate our machine learning and auxiliary channel–based approach because it represents a well-characterized area that has been studied with multiple different approaches [34,57].

A. Scattered light glitches

On December 1, 2019, the LLO detector experienced a high level of scattered light glitches, one of which is illustrated in a time-frequency plot in Fig. 3. Scattering is observed as wide “arches” on the Omegagram, a time-frequency representation of the gravitational-wave channel. These glitches pollute the low-frequency regime (< 100 Hz) of the strain data. This lowest frequency part of the LIGO

sensitivity region is observationally important, as the highest-mass binary black holes observed by LIGO merge in this regime. As a result, LIGO observed an increased trigger rate in the output of compact binary coalescence pipelines, especially for short duration templates [65]. Recognizing, characterizing, and mitigating scattered light glitches thus can significantly increase the confidence in LIGO gravitational-wave detections.

Increased scattered light glitch rates have been observed to correlate with increased microseismic activity at the LIGO detector sites [57,65]. Such an increase occurred at LLO on December 1, 2019, as illustrated in Fig. 4. The “microseism” band refers to the 0.03–0.5 Hz region in the amplitude spectral density curve of the observed ground motion produced by ocean waves. LLO generally experiences increased microseismic activity due to its close proximity to the Gulf of Mexico. The microseismic noise at LLO usually peaks approximately

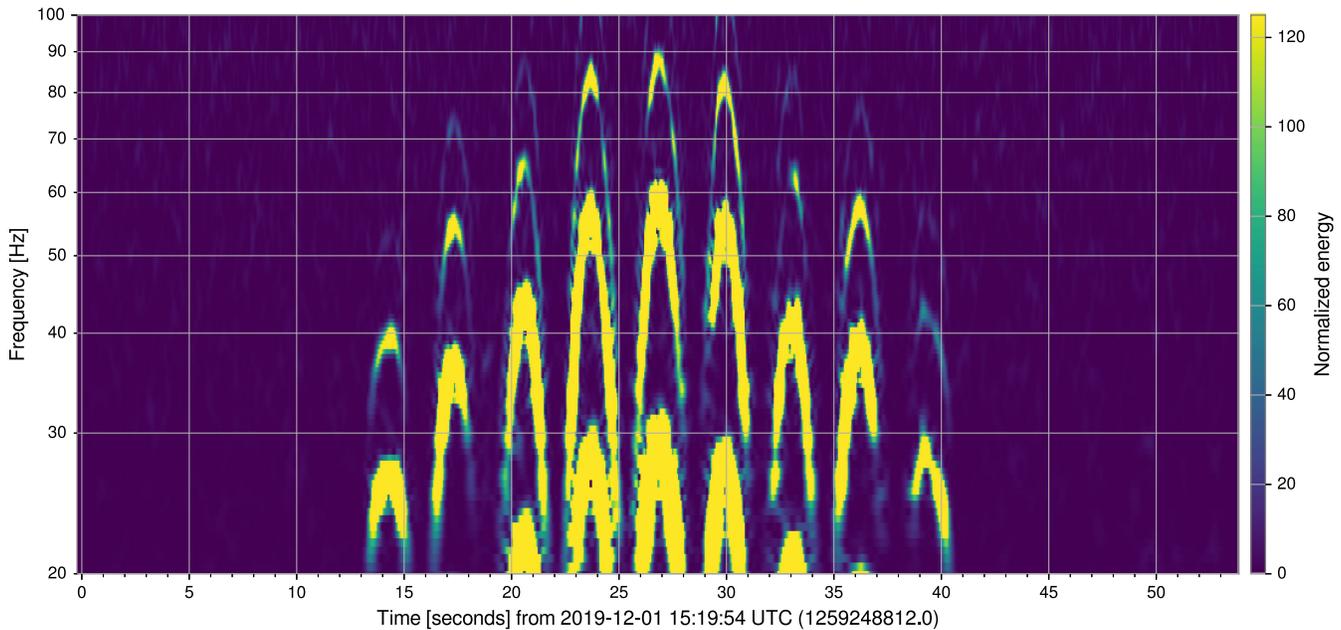


FIG. 3. Omegagram (a Gabor-wavelet time-frequency representation) of the strain data [66,67] of a loud scattered light glitch around GPS time 1,259,248,838.813. Note the characteristic arches with peaks spaced approximately 3–4 seconds apart.

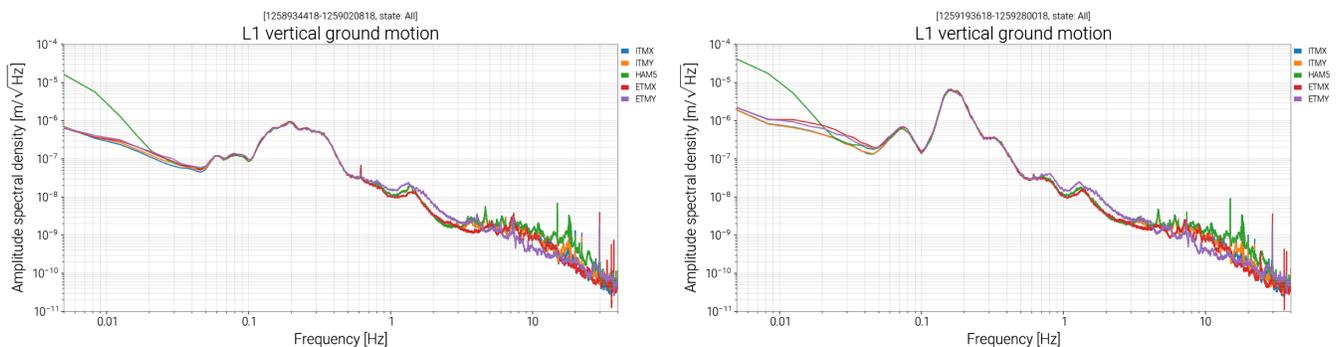


FIG. 4. Microseismic noise on (left) a typical day at LLO and (right) December 1, 2019, showing significantly greater activity in the 0.1–0.3 Hz range.

between 0.1–0.3 Hz. This frequency is below the dominant pendulum frequency of the quadruple suspension (0.45 Hz [57]); therefore, any force applied below this frequency will move the end test mass suspension chain together and allow spurious optical noise to be fed back to the interferometer’s output, giving rise to scattering arches.

B. Data for glitch analysis

We obtained raw auxiliary channel data from the 70,000-second (17 hour, 40 minute) period from GPS times 1,259,197,952 to 1,259,267,952 (December 1, 2019, 1:12:14 to 20:38:54 UTC). During that period, the Omicron glitch detector [55], which monitors the gravitational-wave strain data for events of excess power, recorded a total 48,879 glitches of signal-to-noise ratio (SNR) 5 or

greater. The GravitySpy project [34] selects a subset of high-SNR Omicron glitches that are well suited to morphological classification through a combination of convolutional neural network (CNN)–based machine learning and citizen science. Of the 6,515 glitches classified in total by GravitySpy during that period, 5,499 were classified as “scattered light.” There were several breaks in lock during the above period; we discarded data that fell outside of a lock segment.

We draw training data from the 50,000-second period from GPS times 1,259,197,952 to 1,259,247,952 and validation and test data respectively from the 10,000-second periods from 1,259,247,952 to 1,259,257,952 and 1,259,257,952 to 1,259,267,952. There were a total of 3,555, 924, and 1,016 scattered light glitches during the training, validation, and test periods respectively.

We follow the same procedure as [21,22] to reduce the number of auxiliary channels considered from approximately 250,000 to approximately 40,000 by excluding channels that are constant or only vary in a predictable fashion (e.g., counting time cycles). We also exclude any auxiliary channels known or suspected to be coupled to the gravitational-wave data stream following the procedure of [22]. We include all remaining channels at all sample rates (a total of 39,147). We normalize each channel by computing the mean and standard deviation of the raw channel data over the entire training data period; then we subtract the training mean and divide by the standard deviation for all data in the training, validation, and test periods.

Our positive samples are drawn from points in time identified by GravitySpy [34] as scattered light glitches. As in [21,22], our negative samples are drawn randomly from periods where no glitch was identified by the Omicron glitch detector [55] within two seconds. We select the same number of negative samples as there are positive samples in each dataset.

IV. RESULTS

In this section we describe our findings upon applying the LF model to the scattered light data described in Sec. III. We note that neither the model nor the authors incorporated any specific prior knowledge of the characteristics of scattered light glitches nor the input auxiliary channels; these results were achieved independently and automatically by the method presented here.

The model achieves an overall 97.1% accuracy distinguishing between glitches and glitch-free points from our test dataset. It is able to achieve this level of accuracy after being trained on data from all 39,147 potentially informative auxiliary channels and automatically selecting only 25 of them (Table I) as relevant to predict this type of glitch. For each selected channel, the model also provides a visually interpretable feature for each selected channel, suggesting a signature of behavior in that channel associated with the presence or absence of a scattered light glitch.

These results and their potential interpretations may suggest novel directions for investigation of the origins and causes of scattered light glitches deserving of deeper investigations by instrument science domain experts.

A. Predictive accuracy

The best-performing (by loss) model on the scattered light validation dataset achieved a loss of 0.0933, corresponding to an accuracy of 97.0% (true positive rate 97.4%, true negative rate 96.6%). On the test dataset, it achieved a loss of 0.0866, corresponding to an accuracy of 97.1% (true positive rate 98.4%, true negative rate 95.8%). An ROC curve for the test dataset is shown in Fig. 5.

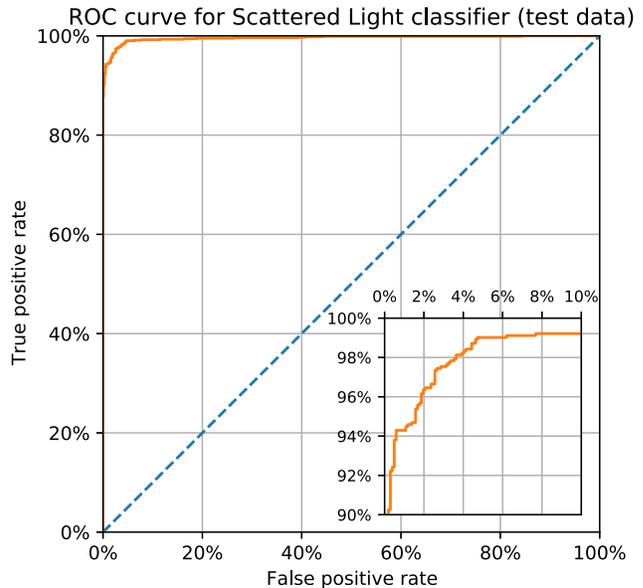


FIG. 5. ROC curve of LF classifier on scattered light test dataset.

Additionally, we manually examined Omegagrams around the 43 instances classified as a glitch by our model but not by GravitySpy or Omicron from our test dataset and found faint but clear signatures of possible scattered light glitches in as many as 30 of them. If we excluded those 30 from the 1,016 samples labeled “glitch-free” in our ground truth, our model’s true negative rate would rise from 95.8% to 98.7% and overall accuracy would rise from 97.1% to 98.6%. We show an example of one such false positive by our model in Fig. 6.

For comparison, we also performed an experiment with the same type of model and with data drawn from the same periods as for the scattered light classifier (as described in Sec. III A) but using all Omicron glitches as the positive class rather than scattered light glitches only. It achieved a maximum validation accuracy of 82.8%, indicating that

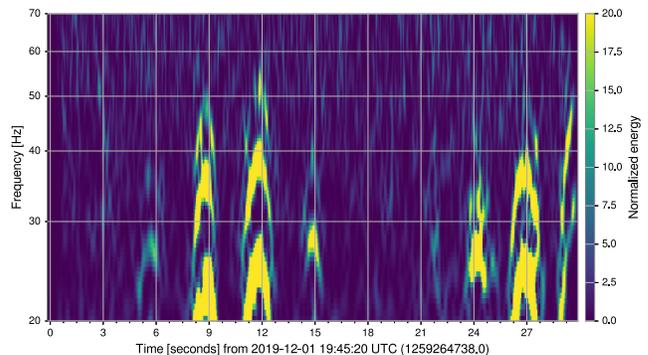


FIG. 6. Omegagram around GPS time 1259264752.9 (14.9 s on the x-axis), which was classified as a glitch by our model but not by GravitySpy or Omicron.

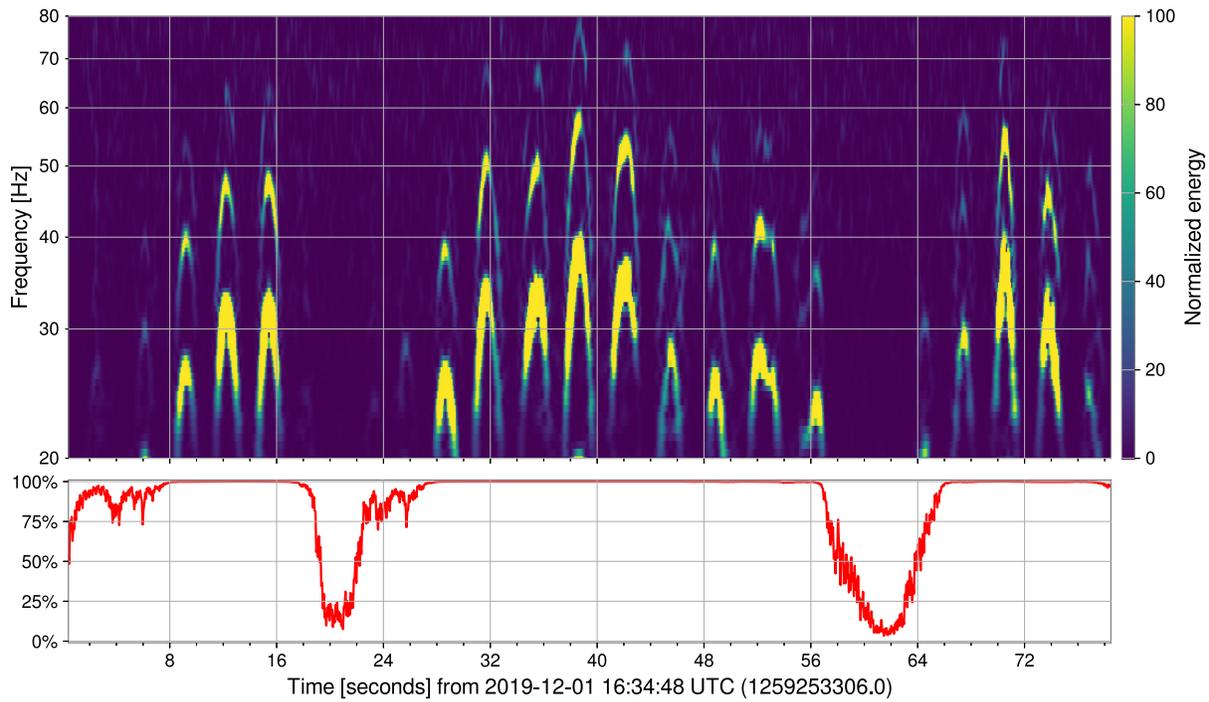


FIG. 7. Omegagram of several scattered light glitches around GPS time 1,259,253,345 (top) and our model’s estimate over time of the probability that there is a glitch (bottom). GravitySpy reports multiple scattered light glitches during this period, including a 19.75-second-long one centered at GPS time 1,259,253,341.438 (approximately 35.4 seconds on the x-axis), a 9.5-second-long one centered at 1,259,253,318.188 (12.2 seconds), and several shorter ones from 1,259,253,370.438 to 1,259,253,382.719 (64.4 to 76.7 seconds).

scattered light glitches are significantly easier to classify than general Omicron glitches.

As an example of how the model’s output varies over time depending on the presence or absence of a glitch, we fed it a continuous segment of auxiliary channel data around several scattered light glitches that occurred around GPS time 1,259,253,345. Figure 7 shows an Omegagram [66,67] of these glitches (top) and our model’s estimate of the probability of a glitch at any given time (bottom). GravitySpy reports a 19.75-second scattered light glitch centered at 1,259,253,341.438 (35 seconds on the x-axis of

Fig. 7), which our model classified as a glitch with very high confidence. GravitySpy also reports several other shorter scattered light glitches around that time, which are visible in the Omegagram and reflected in our model’s output.

We illustrate the Omegagrams of two more scattered light glitches along with the model’s output over the surrounding time period in Figs. 8 and 9. We illustrate the corresponding raw channel value and correlation for selected channels in Figs. 13 and 14 (see Sec. IV D). In Fig. 10, we show the Omegagram of a glitch-free point

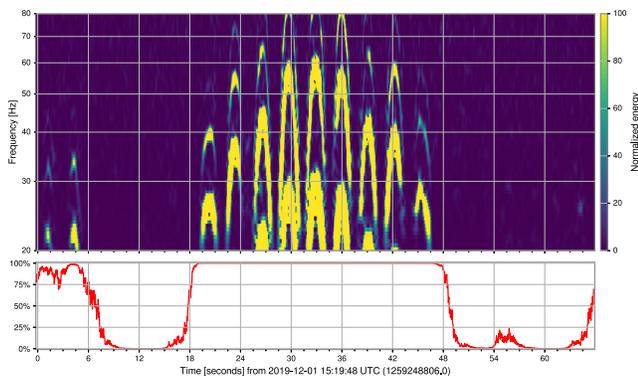


FIG. 8. Omegagram of several scattered light glitches around GPS time 1,259,248,839 (top) and our model’s estimate of the probability that there is a glitch over time (bottom).

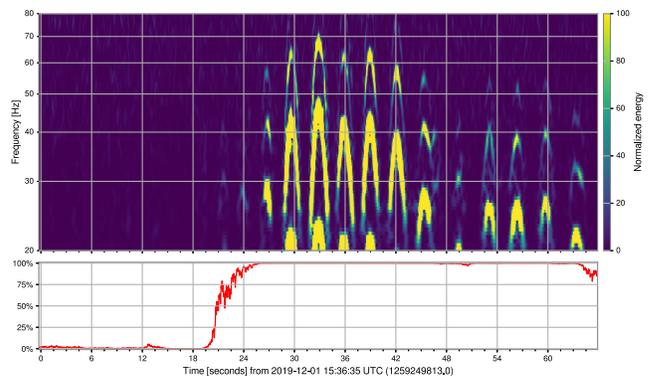


FIG. 9. Omegagram of several scattered light glitches around GPS time 1,259,249,846 (top) and our model’s estimate of the probability that there is a glitch over time (bottom).

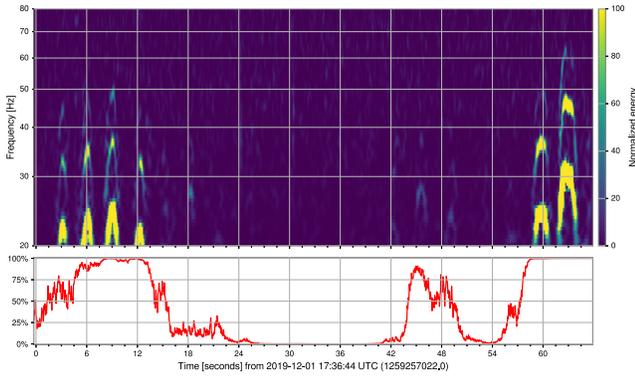


FIG. 10. Omegagram of a glitch-free period of about 30 seconds around GPS time 1,259,257,055 (top) and our model’s estimate of the probability that there is a glitch over time (bottom).

along with the model’s output over the surrounding time period; in Fig. 15, we also show the corresponding raw channel value and correlation for the same channels as shown in Fig. 12.

B. Selected channels

The nonzero channels from the best-performing scattered light model and the magnitudes of their associated learned filters are listed in Table I. Only 25 of the 39,147

channels have associated filters with magnitude greater than zero.

The magnitude of the filters also provides an indication of the relative importance of each channel to the model’s predictions (intuitively, when correlated with a higher-magnitude filter, an input data segment will contribute more heavily to the sum and resulting probability estimate than the same segment correlated with a lower-magnitude filter) [21]. Among those 25, the magnitudes decay rapidly, suggesting the model’s predictions are dominated by the contributions of even fewer channels.

Many of the channels in question are related to LLO’s suspension system (SUS), test masses (ETMX, ITMX), and sensors (OSEM, OPLEV). The details of these subsystems are described in [68–70].

C. Learned filters

Several of the learned filters are illustrated in Fig. 11. We leave detailed interpretive analysis and instrument hardware-related detective work regarding the channels and features themselves to future work. However, we qualitatively observe that many of the channels have a similar rising oscillatory shape. All four of the highest-magnitude channels listed in Table I, the first two of which are shown in the top left and top center of Fig. 1, are nearly identical in shape, likely indicating that they are sensitive to the same or

TABLE I. Channels with nonzero filters and their associated magnitudes from the LF model trained on the scattered light dataset.

Filter magnitude	Channel name	Frequency
0.258	L1:SUS-ETMX_L2_WD_OSEMAC_UL_RMSSON	16 Hz
0.180	L1:SUS-ETMX_L2_WD_OSEMAC_UR_RMSSON	16 Hz
0.117	L1:SUS-ETMX_L2_WD_OSEMAC_LR_RMSSON	16 Hz
0.112	L1:SUS-ETMX_L2_WD_OSEMAC_LL_RMSSON	16 Hz
0.093	L1:SUS-SR2_M1_WD_OSEMAC_T2_RMSSON	16 Hz
0.066	L1:SUS-SR2_M2_RMSIMON_LR_OUT16	16 Hz
0.056	L1:SUS-ETMX_L1_WD_OSEMAC_UL_RMSSON	16 Hz
0.054	L1:SUS-ITMY_L3_OPLEV_BLRMS_S_100M_300M	16 Hz
0.050	L1:SUS-ETMX_R0_WD_OSEMAC_RT_RMSSON	16 Hz
0.046	L1:SUS-ITMX_M0_WD_OSEMAC_SD_RMSSON	16 Hz
0.043	L1:SUS-ITMX_L3_OPLEV_BLRMS_P_1_3	16 Hz
0.038	L1:SUS-ITMY_L3_ISCINF_L_IN1_DQ	16 384 Hz
0.031	L1:IOP-SUS_OMC_WD_OSEM4_RMSOUT	16 Hz
0.031	L1:ISI-BS_ST2_BLRMS_X_300M_1	16 Hz
0.029	L1:SUS-ETMX_L1_WD_OSEMAC_UR_RMSSON	16 Hz
0.022	L1:CAL-PCALY_IRIGB_DQ	16 384 Hz
0.014	L0:FMC-CS_AHU1_FAN1_VS	16 Hz
0.011	L1:SUS-ETMY_PI_DOWNCONV_DC7_SIG_INMON	16 Hz
0.011	L1:SUS-ETMY_PI_DOWNCONV_DC4_SIG_INMON	16 Hz
0.008	L1:HPI-HAM6_BLRMS_X_30_100	16 Hz
0.007	L1:ASC-AS_B_RF72_I_SUM_OUT16	16 Hz
0.006	L1:SUS-ETMX_M0_WD_OSEMAC_F2_RMSSON	16 Hz
0.002	L1:ASC-AS_B_RF72_I_SUM_INMON	16 Hz
0.002	L1:ASC-AS_B_RF72_I_SUM_OUTPUT	16 Hz
0.000	L1:OAF-ST_SEN2ACT_1_1_IN1_DQ	256 Hz

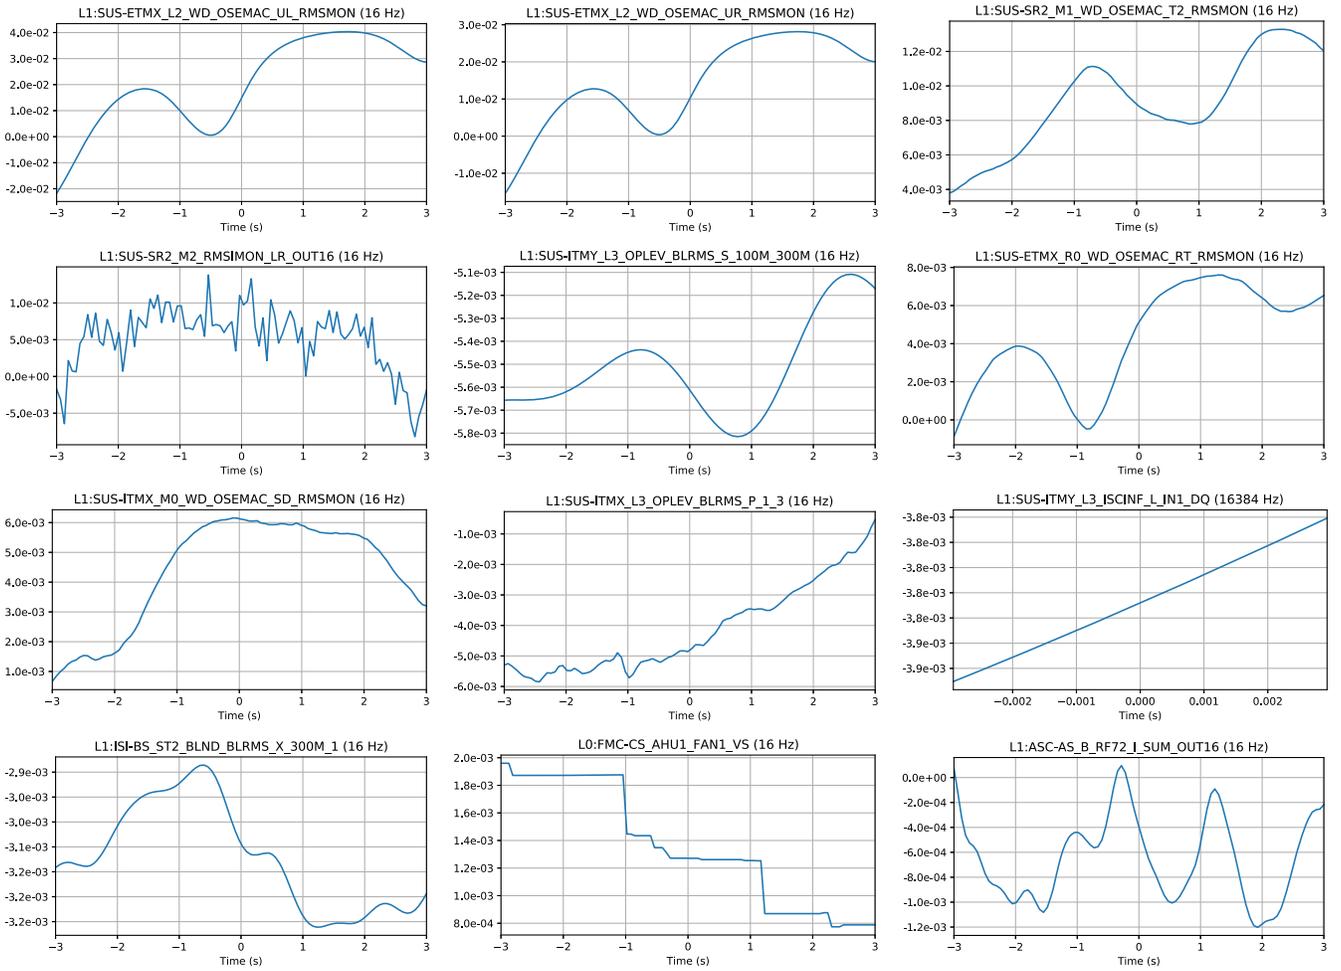


FIG. 11. Twelve of the features learned by the best-performing LF model on the scattered light dataset. The filters shown are ordered by magnitude (see Table I), although we omit some channels for illustration (for example, all of the first four highest-magnitude filters are nearly identical in shape). The x-axes correspond to filter length in seconds; the y-axes show the magnitude of the filter over time. Intuitively, a higher overall magnitude indicates the associated channel is more important to the model’s decisions: when cross-correlated with a higher-magnitude filter, an input data segment will contribute more heavily to the sum and resulting probability estimate than the same segment cross-correlated with a lower-magnitude filter.

similar phenomena; this is also unsurprising given the similarity of their names. We also note that the period of oscillation in the features displaying that pattern is of a similar (3–4 second) duration to the individual arches characteristic of scattered light glitches, perhaps suggesting that these channels are sensitive to the same phenomena that result in the appearance of this type of glitch in the strain.

D. Channel importance for individual glitch classification

Figure 12 shows the normalized value of several auxiliary channels (blue) during the time period containing several scattered light glitches illustrated in Fig. 7 as well as the correlation between the channel and the model’s learned filter (orange) at each time sample during that span. A positive value for the correlation at a given time indicates

the channel is contributing to the model’s classification of the time as containing a glitch (the higher the value the more heavily weighted the contribution), and a negative value indicates the channel is indicating no glitch is present.

We illustrate similar examples with different glitches in Figs. 13 and 14 and an example with a glitch-free point (the same glitch-free point shown in Fig. 10) in Fig. 15.

E. Interpretations

The most significant channels in Table I refer to the third stage (from top) of the quadruple suspension system [71–78] at the X end station (SUS-ETMX_L2). Figures 16 and 17 show the test mass chain and the reaction mass chain behind it. Precision position sensors (so-called OSEMs [73,79–83]) provide the optical sensing and electromagnetic actuating capabilities for the upper three stages. L2 refers to the third (penultimate) stage above the test mass

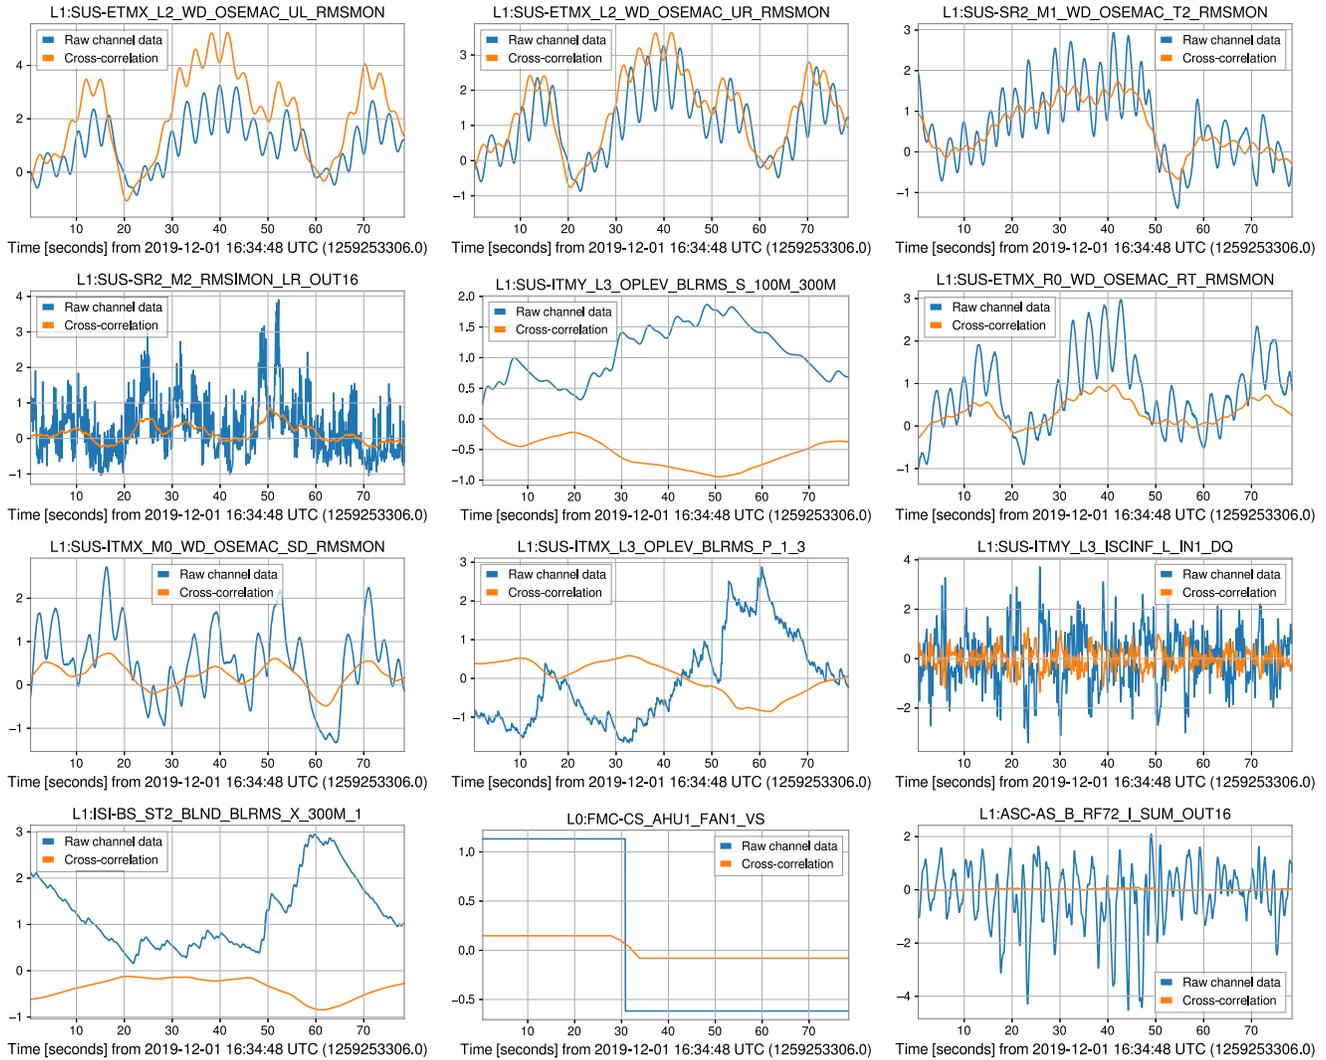


FIG. 12. Raw channel data (blue) from the same time period illustrated in Fig. 7, along with (orange) the cross-correlation of the channel and the associated learned filter (i.e., its contribution to the classifier’s result). Note the oscillations in the “L1:SUS-ETMX” channels with peaks spaced approximately 3.5 seconds apart, similar to the spacing between the peaks of the characteristic scattered light arches in Fig. 7.

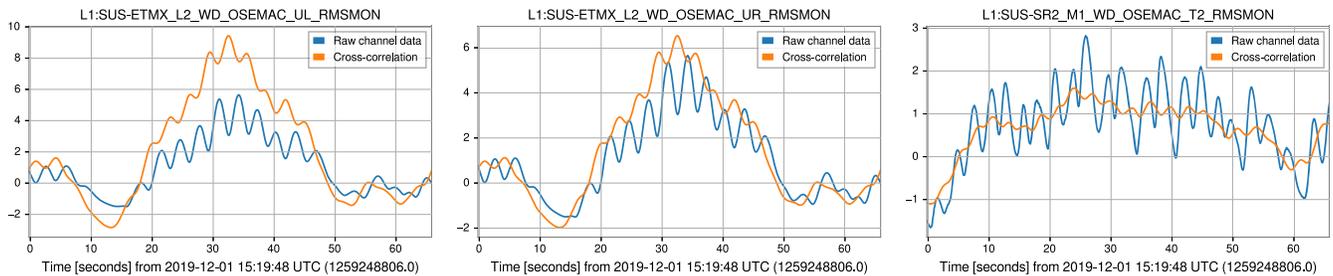


FIG. 13. Raw channel data (blue) from around the scattered light glitch illustrated in Figs. 3 and 8, along with (orange) the output of that channel cross-correlated with the associated learned filter (i.e., its contribution to the classifier’s result). Similarly to Fig. 12, note the oscillations in the channel data with peaks spaced approximately 3.2 seconds apart, similar to the spacing between the peaks of the characteristic scattered light arches in Fig. 3.

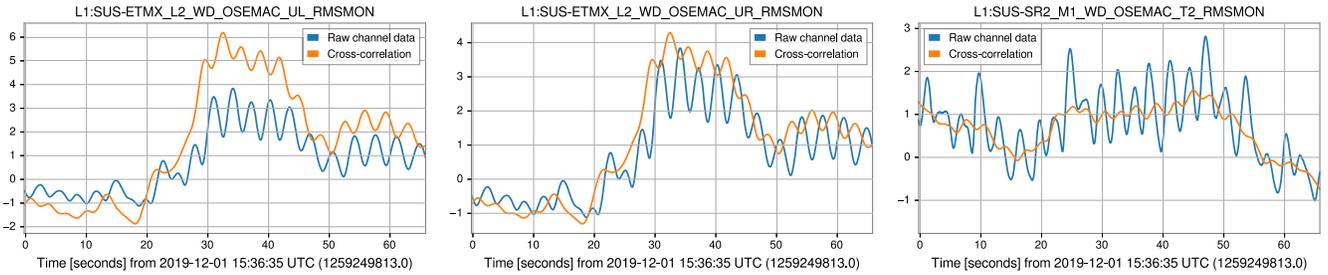


FIG. 14. Raw channel data (blue) from around the scattered light glitch illustrated in Fig. 9, along with (orange) the output of that channel cross-correlated with the associated learned filter (i.e., its contribution to the classifier’s result).

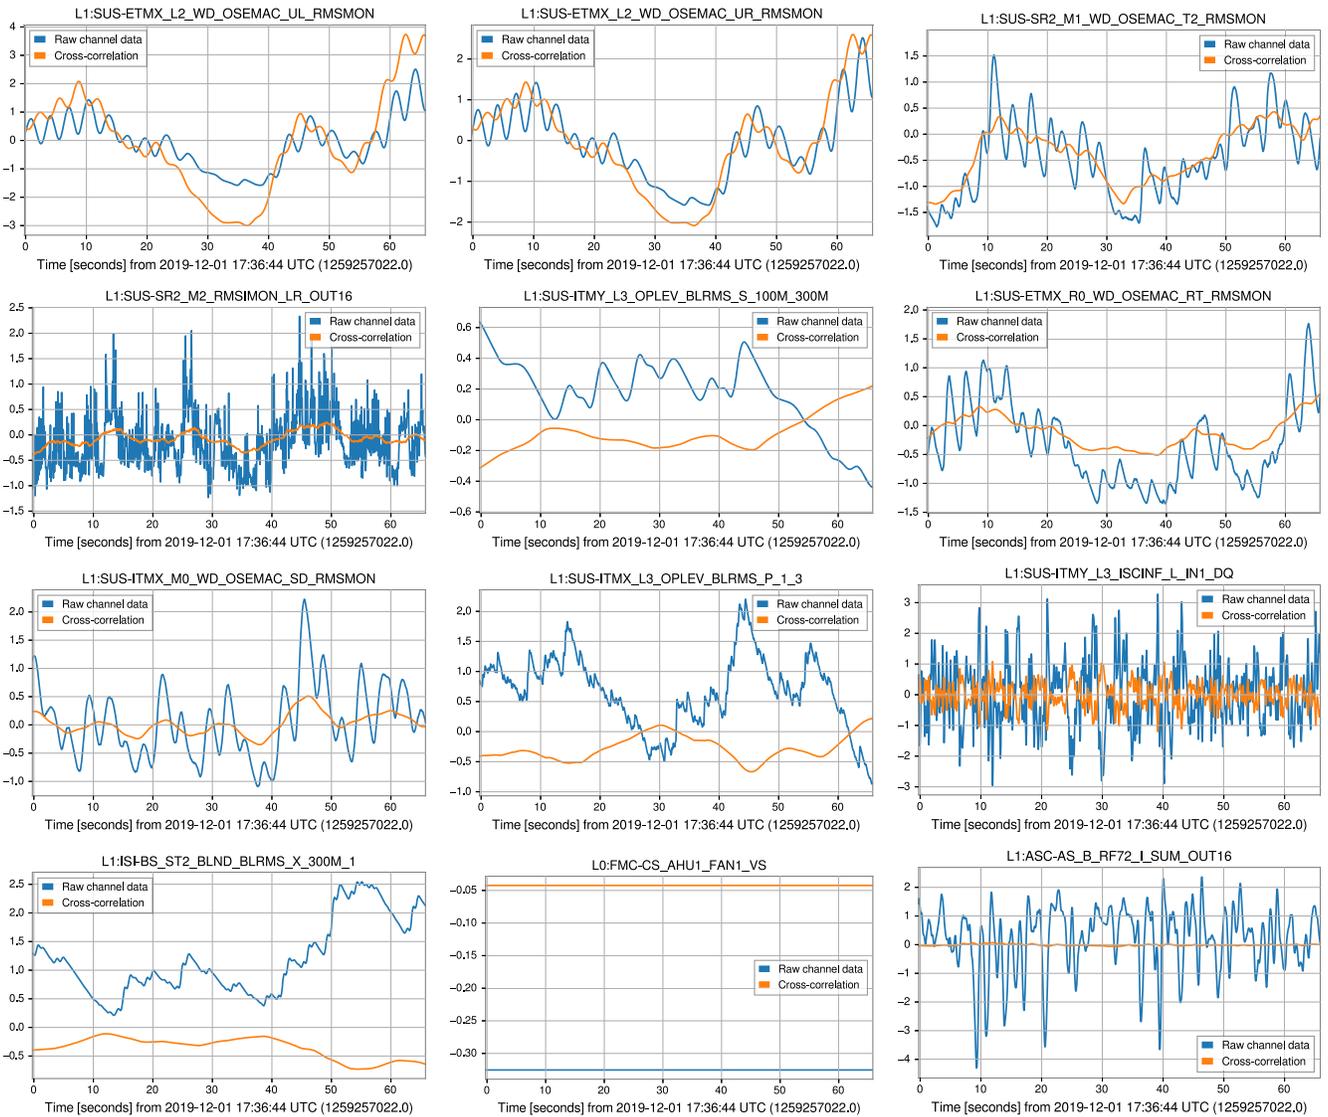


FIG. 15. Raw channel data (blue) from the glitch-free period about 30 seconds long around GPS time 1,259,257,055 illustrated in Fig. 10, along with (orange) the output of that channel cross-correlated with the associated learned filter (i.e., its contribution to the classifier’s result).

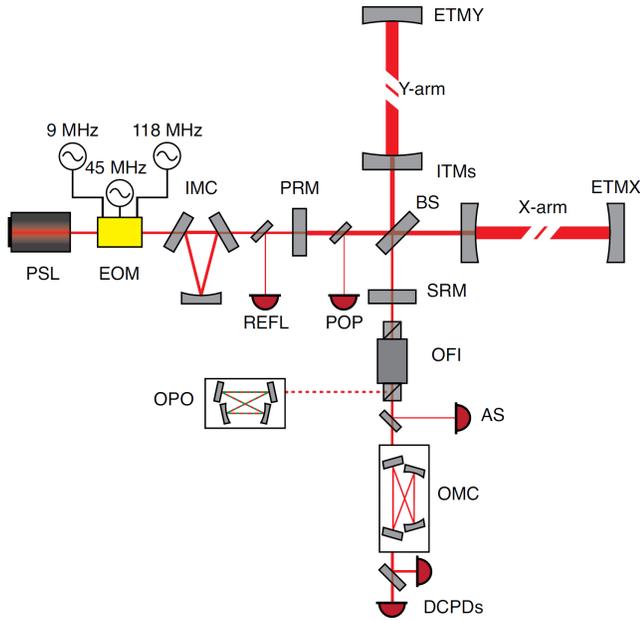


FIG. 16. Diagram of the optical layout the LIGO detectors (from [84]).

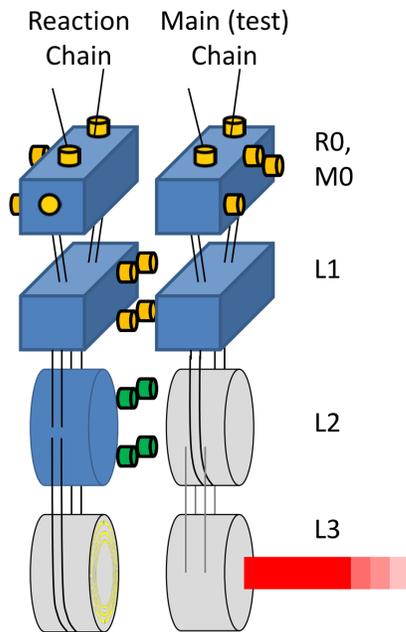


FIG. 17. Diagram of the quadruple suspension system (from [85]).

and the bottom mass of the reaction chain. UL, UR, LL, LR refer to the sensors at the four quadrants of the penultimate mass of the reaction chain. The band-limited RMS motion for these sensors is continuously monitored. For the studied time period, we observe that the channels related to these sensors have the most significant nonzero filters; other channels at SUS-ETMX referring to upper stages of the suspension chain (R0, L1) and other locations around the

detector (e.g., SUS-SR2, ITMX, ITMY) are also indicated, albeit with lower filter magnitudes.

Figure 12 shows the peaks in the oscillations in the L1:SUS-ETMX channels are spaced approximately 3.5 seconds apart, similar to the spacing between the peaks of the characteristic scattered light arches in the corresponding Omegagram (Fig. 7). This suggests that motion indicated by the OSEMs at the end station suspension chain directly gives rise to the observation of scattered light arches in the gravitational-wave strain data channel. Key to taking strain data is controlling the interferometer arms in unison such that the optical cavities in the arms stay resonant and the differential arm length can be continuously sensed. The end test mass chain motion is “locked” to the motion of the other interferometric optics at the detector’s km-scale, and thus its motion relative to other local surfaces in the same chamber can be enhanced, especially during times when high local ground motion is observed.

Figure 4 shows significantly increased (by approximately a factor of five) microseismic activity in the approximately 0.13–0.15 Hz range for the glitch data used in our studies. This frequency can excite motion with a characteristic time of approximately six to eight seconds. Since the spacing of the arches (3–4 seconds as shown in, e.g., Fig. 3) gives half the period of the scattering surface [57], microseism-induced motion is a good candidate for scattering mechanism. This phenomenon was investigated in [57,86,87], and a scattering mechanism involving the bottom-most reaction mass was verified. This ring-shaped mass contains five highly reflective circular gold traces (used as electrostatic drive used for interferometer control). Any incoming stray light will be backscattered toward the test mass with a phase shift; a fraction of that is transmitted into the interferometer arm through the test mass, where it can interfere with the main beam and manifest as an increase in strain amplitude. The path between the reaction mass and end test mass can be traversed multiple times (due to multiple reflections), causing the appearance of harmonics of scattering arches on the Omegagrams (see, for example, Fig. 3). Our study is in agreement with back-scattering from the bottom-most reaction mass being the main source of scattered light glitches in the gravitational-wave data stream. The four most significant channels associated with scattered light glitches directly measure motion of the penultimate reaction mass right above the indicated scattering surface; its motion is naturally connected to that of the scattering surface that can easily couple back to the interferometer arms.

V. CONCLUSION

In this work we have demonstrated the application and capabilities of a machine learning algorithm that employs automatic feature learning to extract relevant information for detecting and analyzing gravitational-wave detector glitches from a vast collection of auxiliary data channels

without the need for manual channel selection or feature engineering. Given tens of thousands of raw time series, the vast majority of which may be irrelevant to the task of glitch analysis, the model automatically selects the few most relevant ones and distills the most pertinent information into several key components that can be further analyzed by detector experts, providing a valuable new tool in their efforts to decrease the incidence of glitches and increase the detectors' sensitivity. In particular, this method provides accurate strain-independent glitch prediction, automatic channel selection, and per-channel behavioral signatures, as well as a straightforward way to investigate which channels are associated with an individual glitch or set of glitches.

The method performs exceptionally well on a real dataset of scattered light glitches from LLO, achieving over 97% accuracy in distinguishing them from glitch-free points. Beyond glitch identification alone, we show how it selects only 25 out of nearly 40,000 channels as relevant to distinguishing glitches from glitch-free times and how it automatically designs relevant features associated with glitches (or the lack thereof) for each of those channels. We also illustrate precisely how the model arrives at its prediction for a given time of interest by visualizing multiple real examples. Unlike some machine learning methods, the method is highly interpretable, allowing detector domain experts to clearly understand how and why it makes a particular individual prediction as well as to probe it for broader understanding of glitch-related behavior over a given time period. We find physical, detector-specific results with meaningful interpretations consistent with previous research on the possible origins of scattered light glitches and their association with microseismic activity near the detector. We also find several instances of scattered light glitches that were missed by Omicron and GravitySpy but successfully identified by our model, as discussed in Sec. IV A and illustrated in Fig. 6, highlighting one of the benefits of having multiple independent, complementary approaches to glitch detection and characterization based on separate data sources.

Previous work [21] has also shown that deeper, more complex neural network-based models can improve performance. In addition to the LF model employed here, [21] also proposed deeper models incorporating nonlinearities, pooling, and other features and showed that they boost performance on a similar LIGO auxiliary channel dataset. However, more complex models are less straightforward to interpret. For the scattered light study presented here (Sec. IV), we found improved performance with the VGG13-BN model of [21]—test accuracy of 98.2% vs 97.1%—but felt that the performance gain was not significant enough to outweigh the tradeoff to interpretability. For other applications or datasets, however, one might prefer to sacrifice some interpretability for improved accuracy. One direction for future work would be to

investigate analytical means of interpreting such models applied to this type of data, especially in situations where performance is more critical or the performance gains are more significant.

Additional directions for future work include evaluating the method on other types of glitches and over longer time periods. Previous work on related models [21,22] has shown that they generally do not perform as well when asked to analyze data drawn from wider time periods, perhaps because they are not able to account well for long-term distribution shifts in the state of the detector. We observed a similar effect in experiments using data drawn from broader timescales. As mentioned in [22], one straightforward method to address this issue would be to periodically (e.g., daily) retrain the model on new labeled data. Another potential avenue would be to identify and test classes of machine learning models that are explicitly designed to handle long-term time dependencies, such as recurrent or long short-term memory neural networks; another would be to retrain or update models like the one employed here online in real time.

We have quantitatively observed here and in previous works that not all glitch morphologies and classes can be equally well predicted based on the auxiliary channels. One intriguing explanation is that information critical to such glitches is not recorded by any current auxiliary channel. Absent such necessary information, it would be difficult or even impossible for any auxiliary channel-based method to detect them. Therefore, we can treat low-accuracy glitch classes as possible clues to additional auxiliary channels that could be added describing previously unmonitored aspects of the detector. Therefore the method can also be useful to help us understand what is missing from the auxiliary monitoring network and channels.

Complex systems of humanity have always been plagued with emergent transient issues that can predict or cause eventual failure. Timely detection, understanding, and repair are paramount to reliable and optimized system operation. Gravitational-wave detectors can serve as excellent test and development environments for generic solutions to unpredictable issues that inevitably emerge in such extremely complex, sensitive systems. The solution outlined above is broadly applicable, enabling automated anomaly discovery, characterization, and human-in-the-loop anomaly elimination. The last step of automated anomaly elimination remains a grand challenge of the field, with the long-term goal of accelerating automated understanding of complex systems and eliminating their problems as they emerge.

ACKNOWLEDGMENTS

We acknowledge computing resources from Columbia University's Shared Research Computing Facility project, which is supported by NIH Research Facility Improvement Grant 1G20RR030893-01, and associated funds from the

New York State Empire State Development, Division of Science Technology and Innovation (NYSTAR) Contract C090171, both awarded April 15, 2010. This material is based upon work and data supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. This research has also made use of data obtained from the Gravitational Wave Open Science Center [88], 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. The authors are grateful for the LIGO Scientific Collaboration review of the paper, with special thanks to Gabriele Vajente. This paper is assigned a LIGO DCC number (LIGO-P2200045). The authors acknowledge the LIGO Collaboration for the production of data used in this study and the LIGO Laboratory for enabling Omicron trigger generation on its computing resources (National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459). The authors are grateful to the authors and maintainers of the Omicron and Omega pipelines, the LIGO Commissioning and Detector Characterization

Teams and LSC domain expert Colleagues whose fundamental work on the LIGO detectors enabled the data used in this paper. The authors would like to thank colleagues of the LIGO Scientific Collaboration and the Virgo Collaboration for their help and useful comments. The authors thank the University of Florida and Columbia University in the City of New York for their generous support. The authors are grateful for the generous support of the National Science Foundation under Grant No. CCF-1740391. I.B. acknowledges the support of the Alfred P. Sloan Foundation and NSF Grants No. PHY-1911796 and No. PHY-2110060.

APPENDIX: TRAINING AND VALIDATION

We follow essentially the same training and validation procedures as [21], except that instead of classifying between all Omicron glitches and glitch-free points, we classify between glitches classified by GravitySpy as “scattered light” and glitch-free points (see Sec. III). We train many models over a grid of hyperparameter settings, evaluate each of them on our validation dataset, and choose the one with the lowest loss. For more detail we refer the reader to [21].

-
- [1] Akutsu *et al.*, Overview of KAGRA: Detector design and construction history, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2021).
- [2] Dooley *et al.*, GEO 600 and the GEO-HF upgrade program: Successes and challenges, *Classical Quantum Gravity* **33**, 075009 (2016).
- [3] J. Aasi, B. Abbott, R. Abbott, T. Abbott, M. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari *et al.*, Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).
- [4] F. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin *et al.*, Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [5] LIGO Scientific and Virgo Collaborations, Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [6] B. P. Abbott *et al.*, GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.* **119**, 161101 (2017).
- [7] B. P. Abbott *et al.*, Multi-messenger Observations of a Binary Neutron Star Merger, *Astrophys. J. Lett.* **848**, L12 (2017).
- [8] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt *et al.*, 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).
- [9] R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, A. Adams, C. Adams, R. X. Adhikari, V. B. Adya, C. Affeldt *et al.* (The LIGO Scientific and the Virgo and the KAGRA Collaborations), 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).
- [10] The LIGO Scientific and the Virgo Collaborations, GWTC-2.1: Deep Extended Catalog of Compact Binary Coalescences Observed by LIGO and Virgo During the First Half of the Third Observing Run, [arXiv:2108.01045](https://arxiv.org/abs/2108.01045).
- [11] The LIGO Scientific, the Virgo, and the KAGRA Collaborations, GWTC-3: Compact binary coalescences observed by ligo and virgo during the second part of the third observing run, [arXiv:2111.03606](https://arxiv.org/abs/2111.03606).
- [12] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, 1-OGC: The first open gravitational-wave catalog of binary mergers from analysis of public advanced LIGO data, *Astrophys. J.* **872**, 195 (2019).
- [13] A. H. Nitz, T. Dent, G. S. Davies, S. Kumar, C. D. Capano, I. Harry, S. Mozzon, L. Nuttall, A. Lundgren, and M. Tápai,

- 2-OGC: Open gravitational-wave catalog of binary mergers from analysis of public advanced LIGO and Virgo data, *Astrophys. J.* **891**, 123 (2020).
- [14] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kasta, M. Schäfer, R. Dhurkunde, and M. Cabero, 3-OGC: Catalog of gravitational waves from compact-binary mergers, *Astrophys. J.* **922**, 76 (2021).
- [15] T. Venumadhav, B. Zackay, J. Roulet, L. Dai, and M. Zaldarriaga, New binary black hole mergers in the second observing run of Advanced LIGO and Advanced Virgo, *Phys. Rev. D* **101**, 083030 (2020).
- [16] B. Zackay, L. Dai, T. Venumadhav, J. Roulet, and M. Zaldarriaga, Detecting gravitational waves with disparate detector responses: Two new binary black hole mergers, *Phys. Rev. D* **104**, 063030 (2021).
- [17] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, Highly spinning and aligned binary black hole merger in the Advanced LIGO first observing run, *Phys. Rev. D* **100**, 023007 (2019).
- [18] V. Gayathri, J. Healy, J. Lange, B. O'Brien, M. Szczepanczyk, I. Bartos, M. Campanelli, S. Klimentko, C. Lousto, and R. O'Shaughnessy, Eccentricity estimate for black hole mergers with numerical relativity simulations, *Nat. Astron.* **6**, 344 (2022).
- [19] 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).
- [20] G. Vajente, Y. Huang, M. Isi, J. C. Driggers, J. S. Kissel, M. J. Szczepańczyk, and S. Vitale, Machine-learning non-stationary noise out of gravitational-wave detectors, *Phys. Rev. D* **101**, 042003 (2020).
- [21] R. E. Colgan, J. Yan, Z. Márka, I. Bartos, S. Márka, and J. N. Wright, Architectural optimization and feature learning for high-dimensional time series datasets, *Phys. Rev. D* **107**, 022009 (2023).
- [22] R. E. Colgan, K. R. Corley, Y. Lau, I. Bartos, J. N. Wright, Z. Márka, and S. Márka, Efficient gravitational-wave glitch identification from environmental data through machine learning, *Phys. Rev. D* **101**, 102003 (2020).
- [23] J. Merritt, B. Farr, R. Hur, B. Edelman, and Z. Doctor, Transient glitch mitigation in Advanced LIGO data with glitschen, *arXiv:2108.12044*.
- [24] D. Davis, LIGO detector characterization in the second and third observing runs, *Classical Quantum Gravity* **38**, 135014 (2021).
- [25] R. Essick, G. Mo, and E. Katsavounidis, A coincidence null test for Poisson-distributed events, *Phys. Rev. D* **103**, 042003 (2021).
- [26] K. Mogushi, Application of a new transient-noise analysis tool for an unmodeled gravitational-wave search pipeline, *Classical Quantum Gravity* **38**, 155004 (2021).
- [27] R. Essick, P. Godwin, C. Hanna, L. Blackburn, and E. Katsavounidis, iDQ: Statistical inference of non-gaussian noise with auxiliary degrees of freedom in gravitational-wave detectors, *Mach. Learn.* **2**, 015004 (2020).
- [28] C. Stachie, T. D. Canton, E. Burns, N. Christensen, R. Hamburg, M. Briggs, J. Broida, A. Goldstein, F. Hayes, T. Littenberg, P. Shawhan, J. Veitch, P. Veres, and C. A. Wilson-Hodge, Search for advanced LIGO single interferometer compact binary coalescence signals in coincidence with Gamma-ray events in Fermi-GBM, *Classical Quantum Gravity* **37**, 175001 (2020).
- [29] D. Davis, L. V. White, and P. R. Saulson, Utilizing aLIGO glitch classifications to validate gravitational-wave candidates, *Classical Quantum Gravity* **37**, 145001 (2020).
- [30] E. Cuoco, Enhancing gravitational-wave science with machine learning, *Mach. Learn. Sci. Tech.* **2**, 011002 (2021).
- [31] R. Gurav, B. Barish, G. Vajente, and E. E. Papalexakis, Unsupervised matrix and tensor factorization for LIGO glitch identification using auxiliary channels, in *AAAI Fall 2020 Symposium on Physics-Guided AI to Accelerate Scientific Discovery* (Association for the Advancement of Artificial Intelligence, 2020).
- [32] M. Razzano and E. Cuoco, Image-based deep learning for classification of noise transients in gravitational wave detectors, *Classical Quantum Gravity* **35**, 095016 (2018).
- [33] M. Cavaglia, K. Staats, and T. Gill, Finding the origin of noise transients in LIGO data with machine learning, *Commun. Comput. Phys.* **25**, 963 (2019).
- [34] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. K. Katsaggelos, S. L. Larson, T. K. Lee, C. Lintott, T. B. Littenberg, A. Lundgren, C. Østerlund, J. R. Smith, L. Trouille, and V. Kalogera, Gravity Spy: Integrating advanced LIGO detector characterization, machine learning, and citizen science, *Classical Quantum Gravity* **34**, 064003 (2017).
- [35] N. Mukund, S. Abraham, S. Kandhasamy, S. Mitra, and N. S. Philip, Transient classification in LIGO data using difference boosting neural network, *Phys. Rev. D* **95**, 104059 (2017).
- [36] G. A. Valdes Sanchez, Data analysis techniques for LIGO detector characterization, Ph.D. thesis, The University of Texas at San Antonio, 2017.
- [37] T. J. Massinger, Detector characterization for Advanced LIGO, Ph.D. thesis, Syracuse University, 2016.
- [38] 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. R. Mullaevy, D. Sigg, J. R. Smith, G. Vajente, A. R. Williamson, and C. C. Wipf, Improving the data quality of Advanced LIGO based on early engineering run results, *Classical Quantum Gravity* **32**, 245005 (2015).
- [39] R. Biswas, L. Blackburn, J. Cao, R. Essick, K. A. Hodge, E. Katsavounidis, K. Kim, Y.-M. Kim, E.-O. Le Bigot, C.-H. Lee, J. J. Oh, S. H. Oh, E. J. Son, Y. Tao, R. Vaulin, and X. Wang, Application of machine learning algorithms to the study of noise artifacts in gravitational-wave data, *Phys. Rev. D* **88**, 062003 (2013).
- [40] D. MacLeod, Improving the sensitivity of searches for gravitational waves from compact binary coalescences, Ph.D. thesis, Cardiff University, United Kingdom, 2013.
- [41] R. Essick, L. Blackburn, and E. Katsavounidis, Optimizing vetoes for gravitational-wave transient searches, *Classical Quantum Gravity* **30**, 155010 (2013).
- [42] J. e. a. Aasi, The characterization of Virgo data and its impact on gravitational-wave searches, *Classical Quantum Gravity* **29**, 155002 (2012).
- [43] J. R. Smith, T. Abbott, E. Hirose, N. Leroy, D. MacLeod, J. McIver, P. Saulson, and P. Shawhan, A hierarchical method

- for vetoing noise transients in gravitational-wave detectors, *Classical Quantum Gravity* **28**, 235005 (2011).
- [44] N. Christensen (LIGO Scientific and Virgo Collaborations), LIGO S6 detector characterization studies, *Classical Quantum Gravity* **27**, 194010 (2010).
- [45] T. Isogai *et al.* (LIGO Scientific and Virgo Collaborations), Used percentage veto for LIGO and virgo binary inspiral searches, *J. Phys. Conf. Ser.* **243**, 012005 (2010).
- [46] S. Mukherjee, R. Obaid, and B. Matkarimov, Classification of glitch waveforms in gravitational wave detector characterization, *J. Phys. Conf. Ser.* **243**, 012006 (2010).
- [47] L. Blackburn, The LSC glitch group: Monitoring noise transients during the fifth LIGO science run, *Classical Quantum Gravity* **25**, 184004 (2008).
- [48] D. Sigg, R. Bork, and J. Zweizig, Detector Characterization and Global Diagnostics System of the Laser Interferometer Gravitational-Wave Observatory (LIGO), in *The Ninth Marcel Grossmann Meeting*, edited by V. G. Gurzadyan, R. T. Jantzen, and R. Ruffini (World Scientific Publishing Co. Pvt. Ltd, Rome, Italy, 2002), pp. 1841–1842.
- [49] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, M. Adamo, C. Adams, T. Adams, P. Addesso *et al.*, Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914, *Classical Quantum Gravity* **33**, 134001 (2016).
- [50] D. Davis and M. Walker, Detector characterization and mitigation of noise in ground-based gravitational-wave interferometers, *Galaxies* **10**, 12 (2022).
- [51] J. D. Merritt, B. Farr, R. Hur, B. Edelman, and Z. Doctor, Transient glitch mitigation in Advanced LIGO data, *Phys. Rev. D* **104**, 102004 (2021).
- [52] P. e. a. Nguyen, Environmental noise in advanced LIGO detectors, *Classical Quantum Gravity* **38**, 145001 (2021).
- [53] L. White (LIGO Team Collaboration), Diamonds in the rough: Characterizing the effect of noise transients on gravitational-wave data, in *APS April Meeting Abstracts, APS Meeting Abstracts* (American Physical Society, Denver, Colorado, 2019), Vol. 2019, p. K01.063.
- [54] D. Davis, Improving the sensitivity of Advanced LIGO through detector characterization, Ph.D. thesis, Syracuse University, USA, 2019.
- [55] F. Robinet, Omicron: An algorithm to detect and characterize transient events in gravitational-wave detectors, <https://tds.ego-gw.it/ql/?c=10651> (2015).
- [56] H. Zou and T. Hastie, Regularization and variable selection via the elastic net, *J. R. Stat. Soc. Ser. B* **67**, 301 (2005).
- [57] S. Soni (The LIGO Scientific Collaboration), Reducing scattered light in LIGO's third observing run, *Classical Quantum Gravity* **38**, 025016 (2021).
- [58] S. Bianchi, A. Longo, G. Valdes, G. González, and W. Plastino, An automated pipeline for scattered light noise characterization, *Classical Quantum Gravity* **39**, 195005 (2022).
- [59] C. D. Austin, Measurements and mitigation of scattered light noise in LIGO, Ph.D. thesis, Louisiana State University, United States, 2020.
- [60] F. Magaña-Sandoval, R. X. Adhikari, V. Frolov, J. Harms, J. Lee, S. Sankar, P. R. Saulson, and J. R. Smith, Large-angle scattered light measurements for quantum-noise filter cavity design studies, *J. Opt. Soc. Am. A* **29**, 1722 (2012).
- [61] The Virgo Collaboration, Characterization of the Virgo seismic environment, *Classical Quantum Gravity* **29**, 025005 (2012).
- [62] J.-Y. Vinet, V. Brisson, S. Braccini, I. Ferrante, L. Pinard, F. Bondu, and E. Tourmié, Scattered light noise in gravitational wave interferometric detectors: A statistical approach, *Phys. Rev. D* **56**, 6085 (1997).
- [63] P. Saha, Noise analysis of a suspended high power Michelson interferometer, Ph.D. thesis, Massachusetts Institute of Technology, 1997.
- [64] T. Accadia, F. Acernese, F. Antonucci, P. Astone, G. Ballardin, F. Barone, M. Barsuglia, T. S. Bauer, M. Beker, A. Belletoile *et al.*, Noise from scattered light in Virgo's second science run data, *Classical Quantum Gravity* **27**, 194011 (2010).
- [65] Y. Asali and Z. Marka, Effect of commissioning break on scattering glitches, LIGO document t2000052.
- [66] J. Rollins, Multimessenger astronomy with low-latency searches for transient gravitational waves, Ph.D. thesis, Columbia University (2011).
- [67] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, Multiresolution techniques for the detection of gravitational-wave bursts, *Classical Quantum Gravity* **21**, S1809 (2004).
- [68] F. Matichard *et al.*, Seismic isolation of Advanced LIGO: Review of strategy, instrumentation and performance, *Classical Quantum Gravity* **32**, 185003 (2015).
- [69] D. Sigg, The advanced LIGO detectors in the era of first discoveries, in *Interferometry XVIII*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9960, edited by K. Creath, J. Burke, and A. Albertazzi Gonçalves (SPIE Optical Engineering + Applications, San Diego, California, 2016), p. 996009.
- [70] J. G. Rollins, Distributed state machine supervision for long-baseline gravitational-wave detectors, *Rev. Sci. Instrum.* **87**, 094502 (2016).
- [71] B. N. Shapiro, R. Adhikari, J. Driggers, J. Kissel, B. Lantz, J. Rollins, and K. Youcef-Toumi, Noise and control decoupling of Advanced LIGO suspensions, *Classical Quantum Gravity* **32**, 015004 (2015).
- [72] S. M. Aston *et al.*, Update on quadruple suspension design for Advanced LIGO, *Classical Quantum Gravity* **29**, 235004 (2012).
- [73] L. Carbone, S. M. Aston, R. M. Cutler, A. Freise, J. Greenhalgh, J. Heefner, D. Hoyland, N. A. Lockerbie, D. Lodhia, N. A. Robertson, C. C. Speake, K. A. Strain, and A. Vecchio, Sensors and actuators for the Advanced LIGO mirror suspensions, *Classical Quantum Gravity* **29**, 115005 (2012).
- [74] A. V. Cumming *et al.*, Design and development of the advanced LIGO monolithic fused silica suspension, *Classical Quantum Gravity* **29**, 035003 (2012).
- [75] N. A. Robertson and E. Majorana, Test mass suspensions, in *Advanced Interferometric Gravitational-Wave Detectors. Volume I: Essentials of Gravitational-Wave Detectors*, edited by D. Reitze *et al.* (World Scientific Publishing, Singapore, 2019), pp. 423–457.
- [76] L. Barsotti, M. Evans, and P. Fritschel, Alignment sensing and control in advanced LIGO, *Classical Quantum Gravity* **27**, 084026 (2010).

- [77] W. Hua, R. Adhikari, D. B. DeBra, J. A. Giaime, G. D. Hammond, C. Hardham, M. Hennessy, J. P. How, B. T. Lantz, M. Macinnis, R. Mittleman, S. Richman, N. A. Robertson, J. Rollins, D. H. Shoemaker, and R. T. Stebbins, Low-frequency active vibration isolation for advanced LIGO, in *Gravitational Wave and Particle Astrophysics Detectors*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 5500, edited by J. Hough and G. H. Sanders (SPIE Astronomical Telescopes + Instrumentation, Glasgow, United Kingdom, 2004), pp. 194–205.
- [78] R. Abbott *et al.*, Seismic isolation enhancements for initial and Advanced LIGO, *Classical Quantum Gravity* **21**, S915 (2004).
- [79] E. Schwartz, Improving the robustness of the advanced LIGO detectors to earthquakes, *Classical Quantum Gravity* **37**, 235007 (2020).
- [80] M. Walker, Effects of transients in LIGO suspensions on searches for gravitational waves, *Rev. Sci. Instrum.* **88**, 124501 (2017).
- [81] H. P. Daveloza, M. Afrin Badhan, M. Diaz, K. Kawabe, P. N. Konverski, M. Landry, and R. L. Savage, Controlling calibration errors in gravitational-wave detectors by precise location of calibration forces, *J. Phys. Conf. Ser.* **363**, 012007 (2012).
- [82] K. A. Strain and B. N. Shapiro, Damping and local control of mirror suspensions for laser interferometric gravitational wave detectors, *Rev. Sci. Instrum.* **83**, 044501 (2012).
- [83] A. Stochino, B. Abbot, Y. Aso, M. Barton, A. Bertolini, V. Boschi, D. Coyne, R. DeSalvo, C. Galli, Y. Huang, A. Ivanov, S. Marka, D. Ottaway, V. Sannibale, C. Vanni, H. Yamamoto, and S. Yoshida, The Seismic Attenuation System (SAS) for the Advanced LIGO gravitational wave interferometric detectors, *Nucl. Instrum. Methods Phys. Res., Sect. A* **598**, 737 (2009).
- [84] A. Buikema, C. Cahillane, G. Mansell, C. Blair, R. Abbott, C. Adams, R. Adhikari, A. Ananyeva, S. Appert, K. Arai *et al.*, Sensitivity and performance of the Advanced LIGO detectors in the third observing run, *Phys. Rev. D* **102**, 062003 (2020).
- [85] B. Shapiro, Overview of Advanced LIGO suspensions, LIGO document g1100866.
- [86] B. Canuel, E. Genin, G. Vajente, and J. Marque, Displacement noise from back scattering and specular reflection of input optics in advanced gravitational wave detectors, *Opt. Express* **21**, 10546 (2013).
- [87] D. J. Ottaway, P. Fritschel, and S. J. Waldman, Impact of upconverted scattered light on advanced interferometric gravitational wave detectors, *Opt. Express* **20**, 8329 (2012).
- [88] Gravitational Wave Open Science Center, <https://www.gwopenscience.org>.