

Detecting and Diagnosing Terrestrial Gravitational-Wave Mimics Through Feature Learning

Robert E. Colgan^{1,2}, Zsuzsa Márka⁵, Jingkai Yan^{2,3}, Imre Bartos⁴, John N. Wright^{2,3}, and Szabolcs Márka⁶

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

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

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

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

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

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

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 from medical to automotive and space systems, enabling automated anomaly discovery and characterization and human-in-the-loop anomaly elimination.

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 ef-

forts for more than a decade, and have included multiple machine learning–based approaches [20–41]. 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, 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 non-astrophysically-sensitive auxiliary channel data as input [42, 43]. Many other glitch detection methods directly analyze the gravitational-wave strain [21, 23, 25, 30–32, 34–39, 44–49]. However, models trained to replicate their results using only auxiliary channels can provide strain-

independent corroboration of their results. 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 [42].

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 [42, 43] 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. In this section, we describe one such model initially proposed in [43], 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 [50] and GravitySpy [21] 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 not astrophysically sensitive, it is unlikely that these methods would inadvertently identify a true gravitational wave as a glitch. [Section 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 [42, 43]) 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. [Section 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. [Section 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. [Section IV D]

For the experiments presented here, we employ a slightly modified version of the LF model introduced in

[43]. The model is a flat convolutional binary classifier parameterized by a learned filter $\mathbf{f}_p \in \mathbb{R}^T$ 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

To make a prediction at time t , the model takes as input the normalized data (centered at t) for all channels $\mathbf{x}_p[t] \in \mathbb{R}^T$ of each sample rate s_p ; cross-correlates the data with the corresponding filter \mathbf{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,

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

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

As in [43], we employ a sparsifying regularizer—the elastic net [51]—during training to encourage most of the filters to be 0, i.e. $\|\mathbf{f}_p\| = 0$ for most p , because sparsity was found to be essential to good performance as well as beneficial to interpretability [42, 43].

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) [43].

In Figs. 1 and 2, we illustrate visually how the model arrives at its predictions given a potential glitch time t . Fig. 1 shows two examples of learned filters \mathbf{f}_p . Fig. 2 illustrates two channels’ data $\mathbf{x}_p[t]$ (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 section we demonstrate its application to a particularly troublesome glitch type known as “scattered light” [54–60] 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 [21, 54].

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 ($<100\text{Hz}$) 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 [61]. 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 [54, 61]. 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 between 0.1–0.3 Hz. This frequency is below the dominant pendulum frequency of the quadrupole suspension (0.45Hz [54]); therefore, any force applied below this frequency will move the end test mass suspension chain together and allow spurious opti-

¹ We set the filter length T to 96 samples for all channels based on the finding of [43] 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 3/8 seconds of data for a 256-Hz channel, for example.

² The minor modification we make to the LF model of [43] 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. ([43] used only 16-Hz channels in its experiments.)

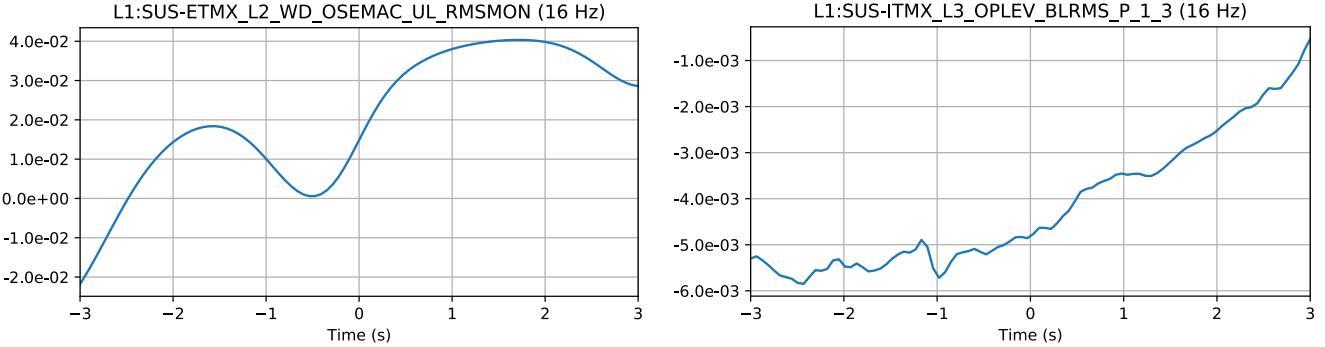


FIG. 1. An example of two learned filters in a trained LF model corresponding to two of the channels found by the model to contain behavior associated with scattered light glitches. See Sec. IV for more details about what the channels in question measure and what behaviors the model has learned to be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

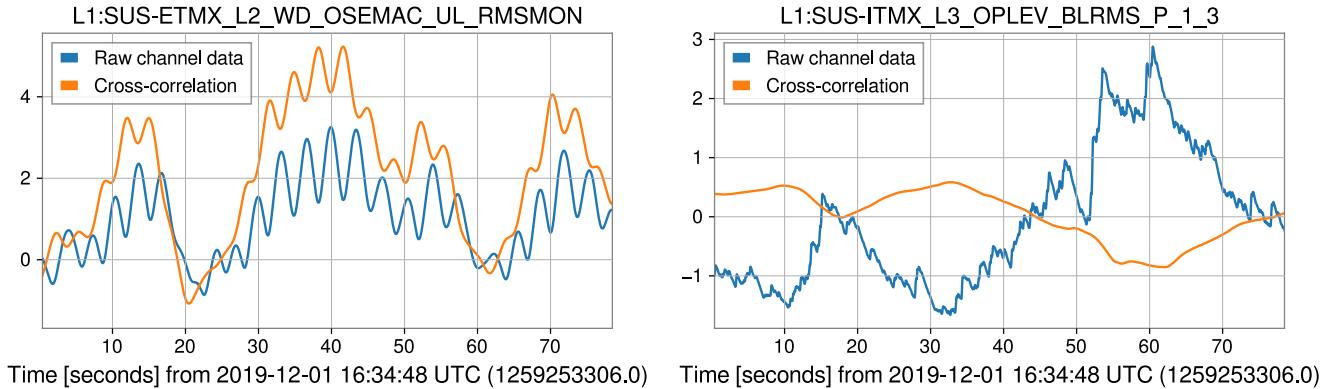


FIG. 2. Blue: normalized channel data from the two channels whose corresponding filters are illustrated in Fig. 1 over time during several scattered light glitches, including a 19.75-second-long one centered at GPS time 1,259,253,341.438 (approximately 35.4 seconds on the x-axis), a 9.5-second-long one 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. Fig. 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.

cal 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 [50], 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 [21] selects a subset of high-SNR Omicron glitches that are well-suited to morphological classification through a combi-

nation 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.”³ 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

³ Of the 6,515 total GravitySpy glitches, all but 297 were also recorded as glitches by Omicron.

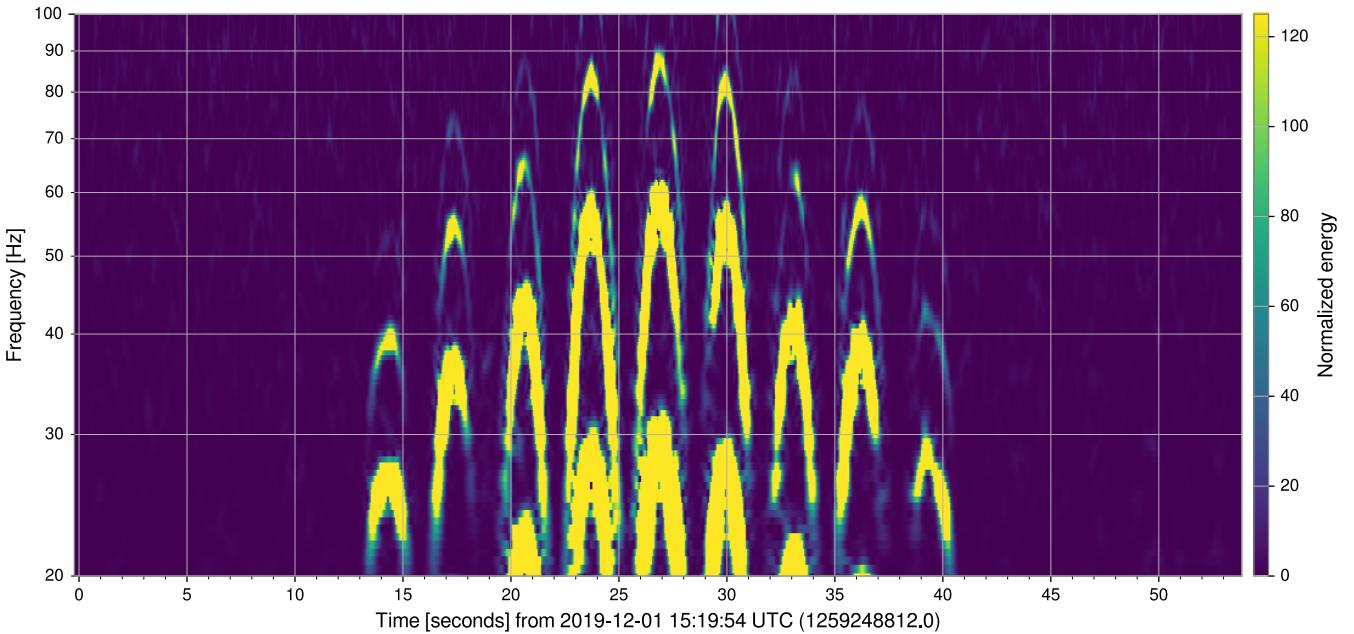


FIG. 3. Omegagram (a Gabor-wavelet time-frequency representation) of the strain data [52][53]) of a loud scattered light glitch around GPS time 1,259,248,838.813. Note the characteristic arches with peaks spaced approximately 3-4 seconds apart.

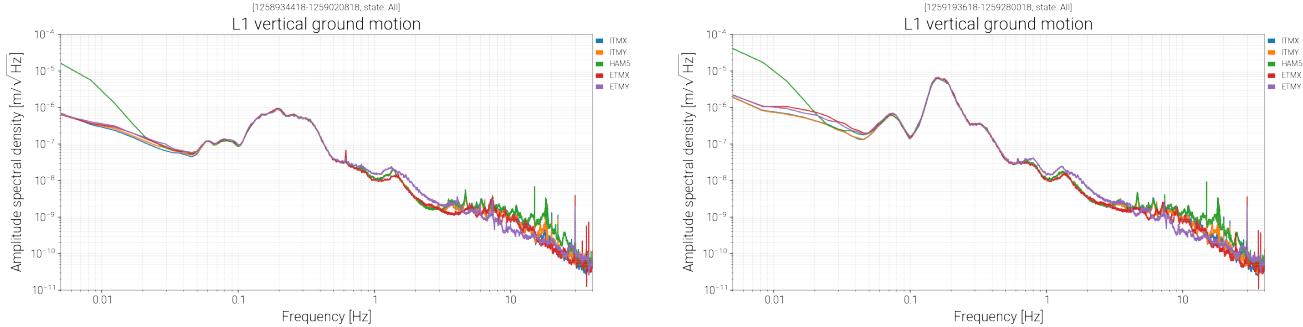


FIG. 4. Microseismic noise on (left) a typical day at LLO and (right) December 1, 2019, showing significantly greater activity in the 0.1-0.3 Hz range.

training, validation, and test periods respectively.

We follow the same procedure as [42] and [43] 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 [42]. 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 [21] as scattered light glitches.

As in [42] and [43], our negative samples are drawn randomly from periods where no glitch was identified by the Omicron glitch detector [50] 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 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 distin-

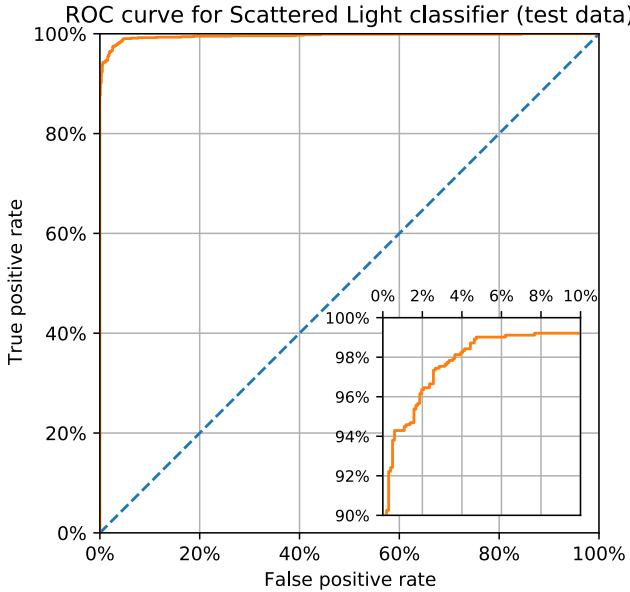


FIG. 5. ROC curve of LF classifier on scattered light test dataset.

guishing 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.

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

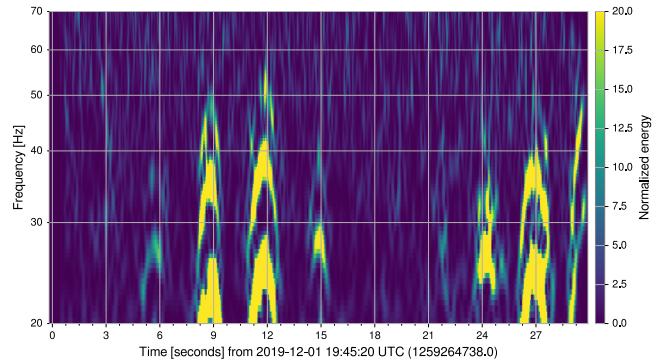


FIG. 6. Omegagram around GPS time 1259264752.9 (14.9 s on the x-axis), which was classified as a glitch by our model but not by GravitySpy or Omicron.

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 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. Fig. 7 shows an Omegagram [52, 53] 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 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.

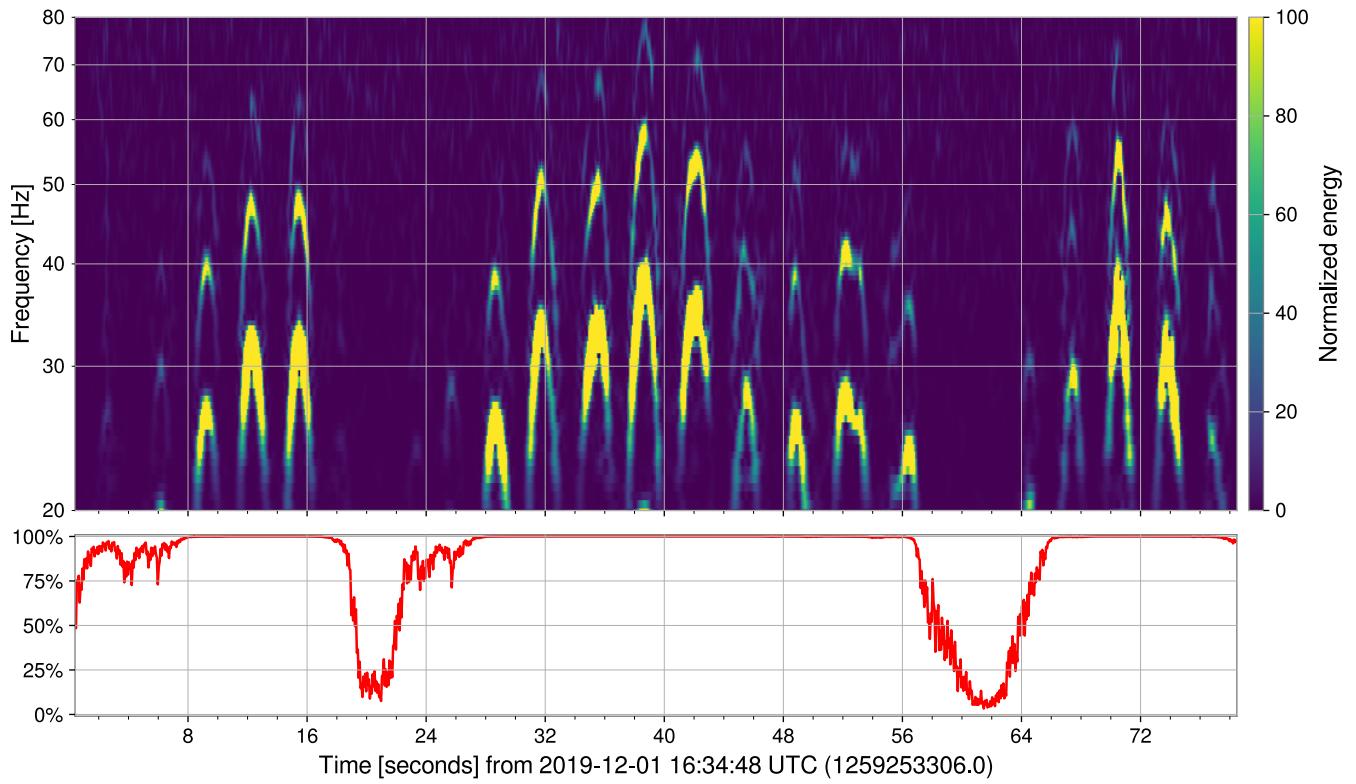


FIG. 7. Omegagram of several scattered light glitches around GPS time 1,259,253,345 (top) and our model's estimate over time of the probability that there is a glitch (bottom). GravitySpy reports multiple scattered light glitches during this period, including a 19.75-second-long one centered at GPS time 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).

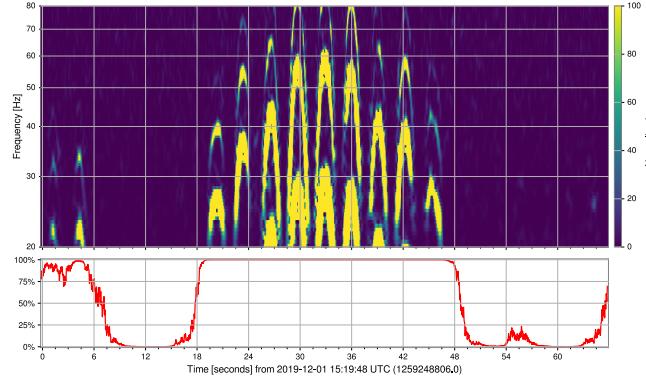


FIG. 8. Omegagram of several scattered light glitches around GPS time 1,259,248,839 (top) and our model's estimate of the probability that there is a glitch over time (bottom).

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

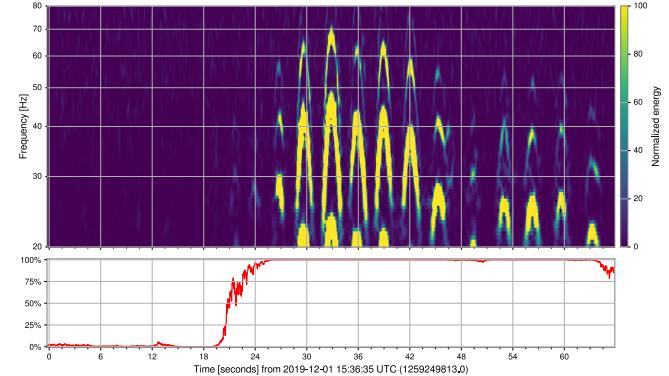


FIG. 9. Omegagram of several scattered light glitches around GPS time 1,259,249,846 (top) and our model's estimate of the probability that there is a glitch over time (bottom).

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 probabil-

Filter magnitude	Channel name	Frequency
0.258	L1:SUS-ETMX_L2_WD_OSEMAC_UL.RMSMON	16 Hz
0.180	L1:SUS-ETMX_L2_WD_OSEMAC_UR.RMSMON	16 Hz
0.117	L1:SUS-ETMX_L2_WD_OSEMAC_LR.RMSMON	16 Hz
0.112	L1:SUS-ETMX_L2_WD_OSEMAC_LL.RMSMON	16 Hz
0.093	L1:SUS-SR2_M1_WD_OSEMAC_T2.RMSMON	16 Hz
0.066	L1:SUS-SR2_M2_RMSIMON_LR.OUT16	16 Hz
0.056	L1:SUS-ETMX_L1_WD_OSEMAC_UL.RMSMON	16 Hz
0.054	L1:SUS-ITMY_L3_OPLEV_BLRMS_S_100M_300M	16 Hz
0.050	L1:SUS-ETMX_R0_WD_OSEMAC_RT.RMSMON	16 Hz
0.046	L1:SUS-ITMX_M0_WD_OSEMAC_SD.RMSMON	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	16384 Hz
0.031	L1:IOP-SUS_OMC_WD_OSEM4.RMSOUT	16 Hz
0.031	L1:ISI-BS_ST2_BLND_BLRMS_X_300M_1	16 Hz
0.029	L1:SUS-ETMX_L1_WD_OSEMAC_UR.RMSMON	16 Hz
0.022	L1:CAL-PCALY_IRIGB_DQ	16384 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.RMSMON	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-STS_SEN2ACT_1_1_IN1.DQ	256 Hz

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

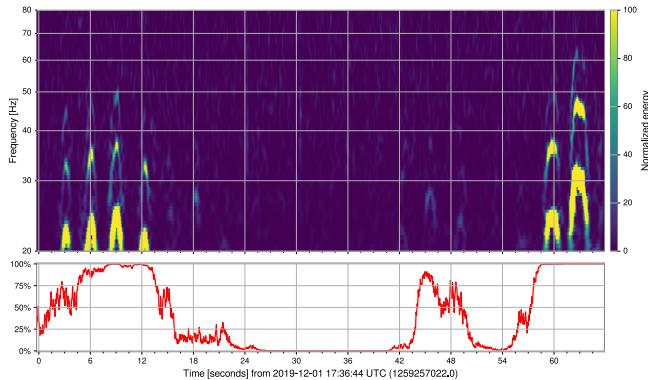


FIG. 10. Omegagram of a glitch-free period of about 30 seconds around GPS time 1,259,257,055 (top) and our model's estimate of the probability that there is a glitch over time (bottom).

ity estimate than the same segment correlated with a lower-magnitude filter) [43]. 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 [62–64].

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 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

Fig. 12 shows the normalized value of several auxiliary channels during that span (blue) as well as the correlation between the channel and the model's learned filter (orange) at each time sample during that span. A pos-

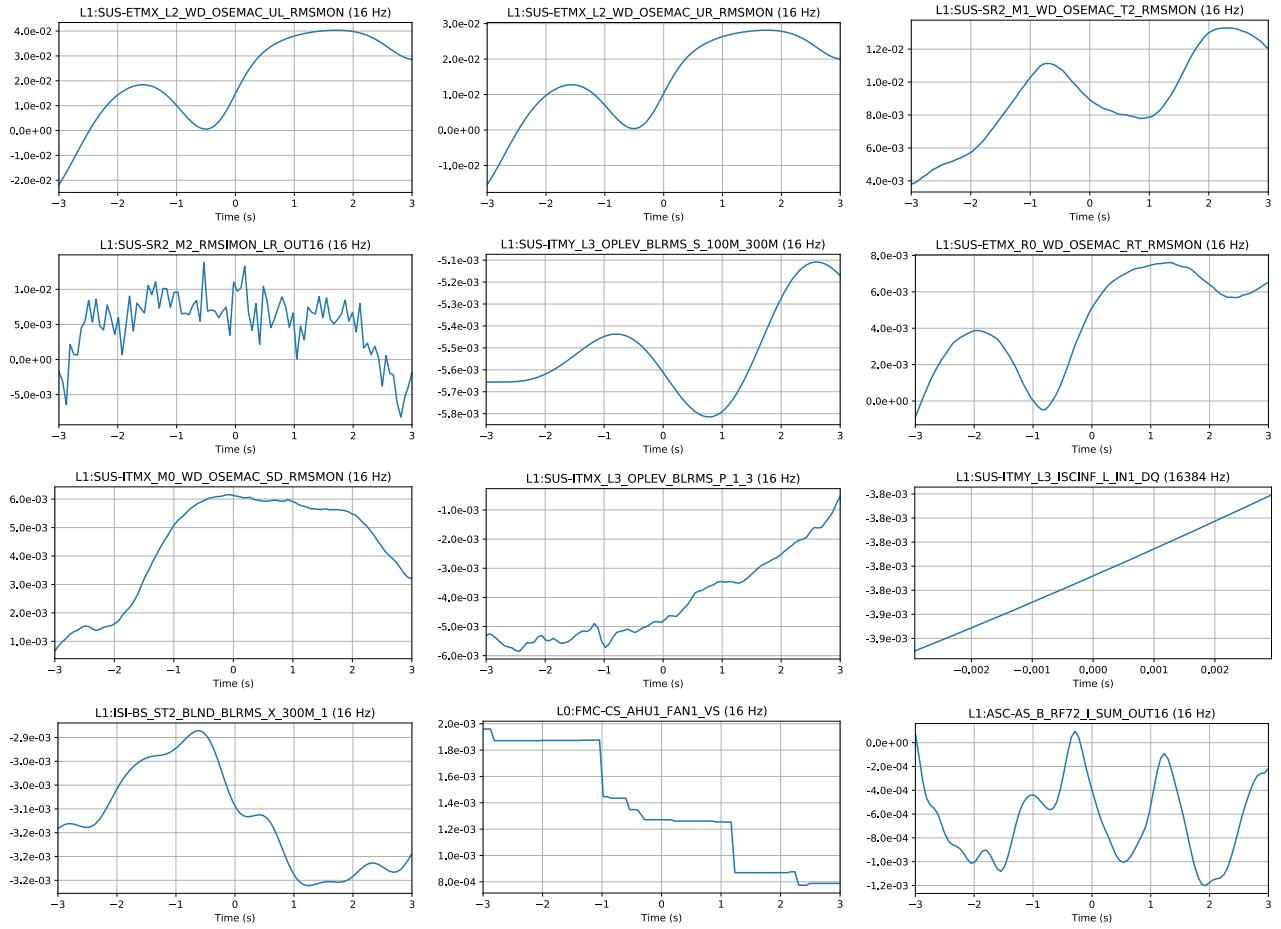


FIG. 11. Twelve of the features learned by the best-performing LF model on the scattered light dataset. The filters shown are ordered by magnitude (see Table I), although we omit some channels for illustration (for example, all of the first four highest-magnitude filters are nearly identical in shape). 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.

itive 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 quadrupole suspension system [67–74] at the X end station (SUS-ETMX.L2). Figs. 16 and 17 show the test mass chain and the reaction mass chain behind it. Precision position sensors (so-called OSEMs [69, 75–79]) provide the optical sensing

and electromagnetic actuating capabilities for the upper three stages. L2 refers to the third (penultimate) stage above the test mass 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. We observe that the channels related to these sensors have the most significant non-zero filters; other channels at SUS-ETMX referring to upper stages of the suspension chain (R0, L1) are also indicated.

Fig. 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 in-

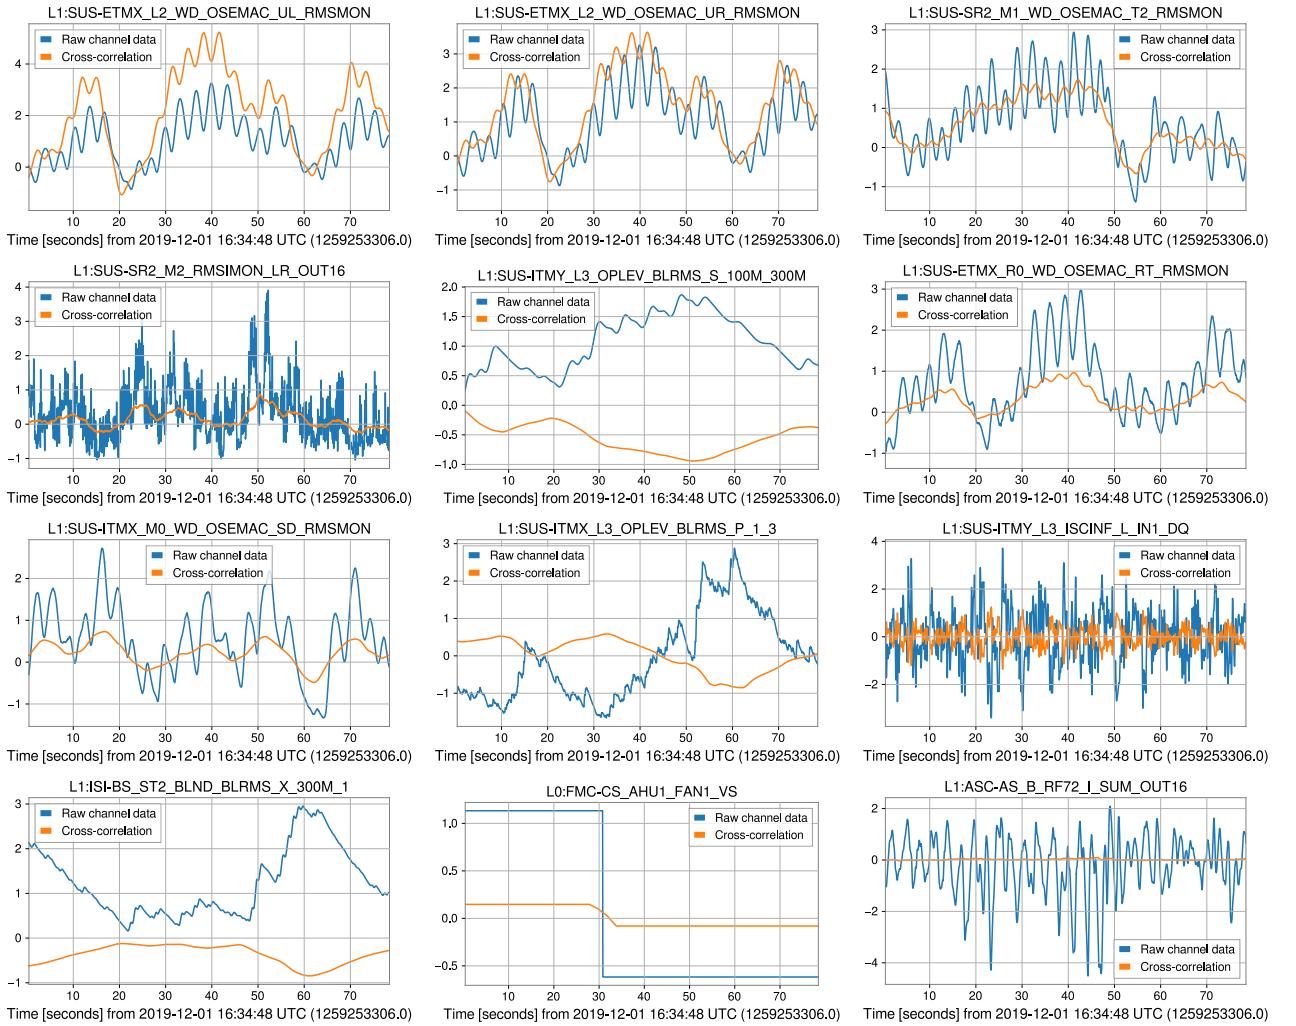


FIG. 12. Raw channel data (blue) from the same time period illustrated in Fig. 7, along with (orange) the cross-correlation of the channel and the associated learned filter (i.e., its contribution to the classifier’s result). Note the oscillations in the “L1:SUS-ETMX” channels with peaks spaced approximately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

terferometer 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.

Fig. 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 phenomenon was investigated in [54, 80, 81], 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 back-scattered towards 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.

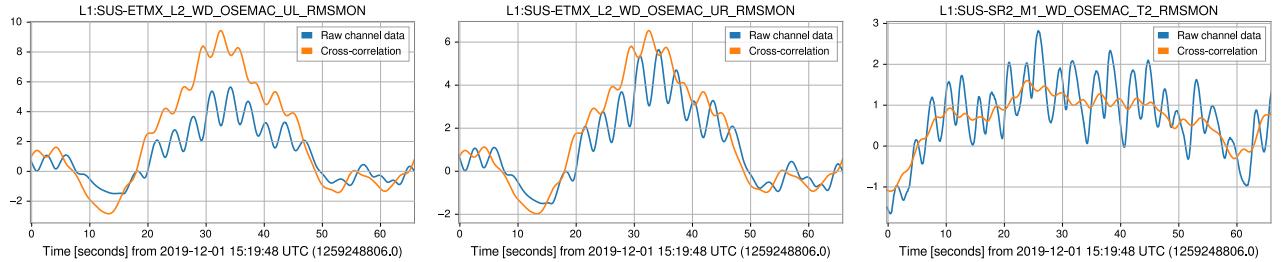


FIG. 13. Raw channel data (blue) from around the scattered light glitch illustrated in Figs. 3 and 8, along with (orange) the output of that channel cross-correlated with the associated learned filter (i.e., its contribution to the classifier’s result). Similarly to Fig. 12, note the oscillations in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

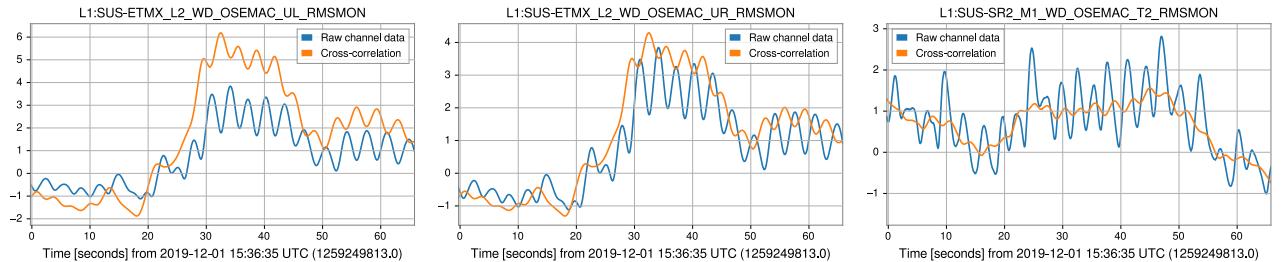


FIG. 14. Raw channel data (blue) from around the scattered light glitch illustrated in Fig. 9, along with (orange) the output of that channel cross-correlated with the associated learned filter (i.e., its contribution to the classifier’s result).

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.

Previous work [43] has also shown that deeper, more complex neural network-based models can improve performance. In addition to the LF model employed here, [43] 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 [43]—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.

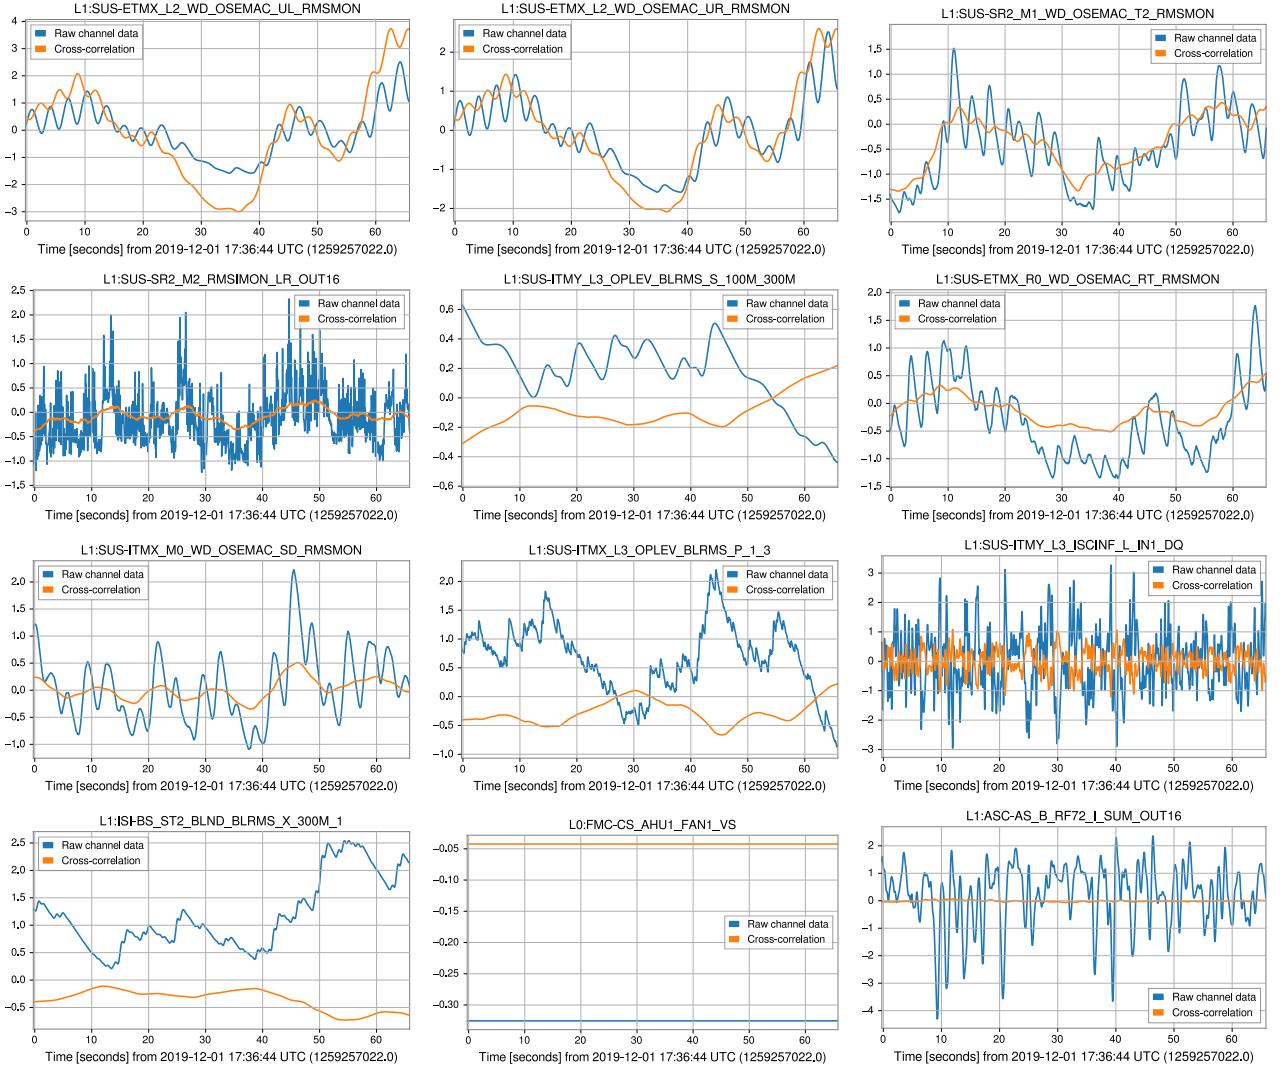


FIG. 15. Raw channel data (blue) from the glitch-free period about 30 seconds long around GPS time 1,259,257,055 illustrated in Fig. 10, along with (orange) the output of that channel cross-correlated with the associated learned filter (i.e., its contribution to the classifier's result).

Additional directions for future work include evaluating the method on other types of glitches and over longer time periods. Previous work on related models 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. One potential avenue to mitigate this 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

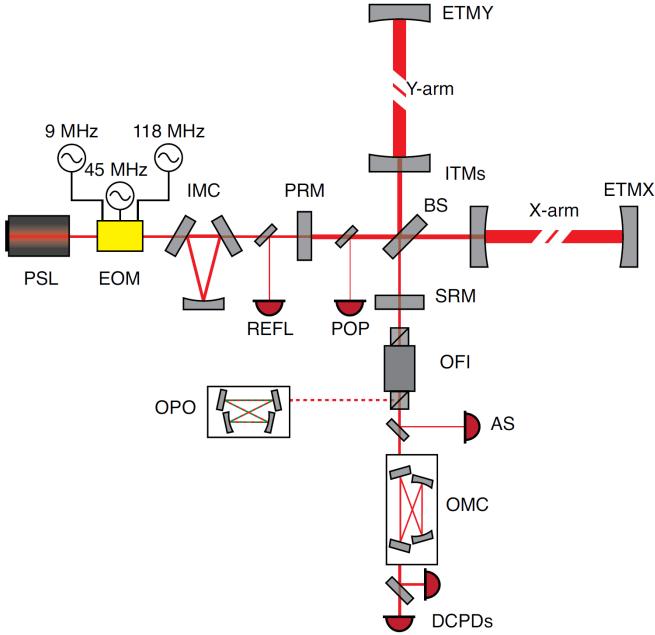


FIG. 16. Diagram of the optical layout the LIGO detectors (from [65]).

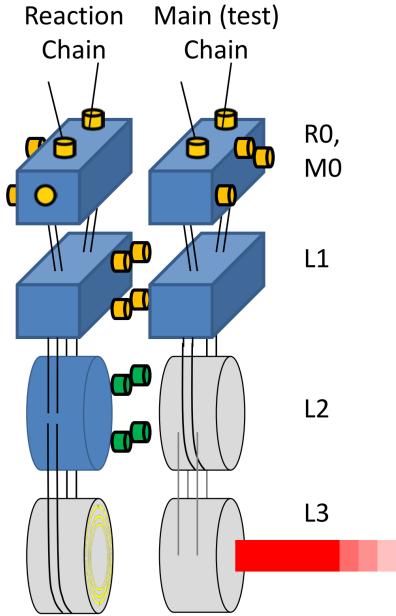


FIG. 17. Diagram of the quadruple suspension system (from [66]).

inevitably emerge in such extremely complex, sensitive systems. The solution outlined above is broadly applicable from medical to automotive and space systems, 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 (<https://www.gw-openscience.org>), a service of LIGO Laboratory, the LIGO Scientific Collaboration and the Virgo Collaboration. LIGO is funded by the U.S. National Science Foundation. Virgo is funded by the French Centre National de Recherche Scientifique (CNRS), the Italian Istituto Nazionale della Fisica Nucleare (INFN) and the Dutch Nikhef, with contributions by Polish and Hungarian institutes.

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 PHY-0757058 and 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 CCF-1740391. I.B. acknowledges the support of the Alfred P. Sloan Foundation and NSF grants PHY-1911796 and PHY-2110060.

Appendix A: Training and Validation

We follow essentially the same training and validation procedures as [43], 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 [43].

[1] Akutsu+, *Progress of Theoretical and Experimental Physics* **2021**, 05A101 (2021), arXiv:2005.05574 [physics.ins-det].

[2] Dooley+, *Classical and Quantum Gravity* **33**, 075009 (2016), arXiv:1510.00317 [physics.ins-det].

[3] J. Aasi, B. Abbott, R. Abbott, T. Abbott, M. Abernathy, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, et al., *Classical and Quantum Gravity* **32**, 074001 (2015), arXiv:1411.4547 [gr-qc].

[4] F. a. Acernese, M. Agathos, K. Agatsuma, D. Aisa, N. Allemandou, A. Allocca, J. Amarni, P. Astone, G. Balestri, G. Ballardin, et al., *Classical and Quantum Gravity* **32**, 024001 (2015), arXiv:1408.3978 [gr-qc].

[5] LIGO Scientific Collaboration and Virgo Collaboration, *Phys. Rev. Lett.* **116**, 061102 (2016), arXiv:1602.03837 [gr-qc].

[6] B. P. Abbott and et al., *Phys. Rev. Lett.* **119**, 161101 (2017).

[7] B. P. Abbott and et al., *Astrophysical Journal Letters* **848**, L12 (2017), arXiv:1710.05833 [astro-ph.HE].

[8] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, and et al., *Physical Review X* **9** (2019), 10.1103/physrevx.9.031040.

[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 Collaboration and the Virgo Collaboration and the KAGRA Collaboration), arXiv e-prints , arXiv:2010.14527 (2020), arXiv:2010.14527 [gr-qc].

[10] The LIGO Scientific Collaboration and the Virgo Collaboration, arXiv e-prints , arXiv:2108.01045 (2021), arXiv:2108.01045 [gr-qc].

[11] The LIGO Scientific Collaboration, the Virgo Collaboration, and the KAGRA Collaboration, arXiv e-prints , arXiv:2111.03606 (2021), arXiv:2111.03606 [gr-qc].

[12] A. H. Nitz, C. Capano, A. B. Nielsen, S. Reyes, R. White, D. A. Brown, and B. Krishnan, *Astrophys. J.* **872**, 195 (2019), arXiv:1811.01921 [gr-qc].

[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, *Astrophys. J.* **891**, 123 (2020), arXiv:1910.05331 [astro-ph.HE].

[14] A. H. Nitz, C. D. Capano, S. Kumar, Y.-F. Wang, S. Kastha, M. Schäfer, R. Dhurkunde, and M. Cabero, arXiv e-prints , arXiv:2105.09151 (2021), arXiv:2105.09151 [astro-ph.HE].

[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,” (2019), arXiv:1904.07214 [astro-ph.HE].

[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,” (2019), arXiv:1910.09528 [astro-ph.HE].

[17] B. Zackay, T. Venumadhav, L. Dai, J. Roulet, and M. Zaldarriaga, *Physical Review D* **100** (2019), 10.1103/physrevd.100.023007.

[18] V. Gayathri, J. Healy, J. Lange, B. O’Brien, M. Szczepanczyk, I. Bartos, M. Campanelli, S. Klimenko, C. Lousto, and R. O’Shaughnessy, (2020), arXiv:2009.05461 [astro-ph.HE].

[19] D. Davis, T. Massinger, A. Lundgren, J. C. Driggers, A. L. Urban, and L. Nuttall, *Classical and Quantum Gravity* **36**, 055011 (2019), arXiv:1809.05348 [astro-ph.IM].

[20] G. Vajente, Y. Huang, M. Isi, J. C. Driggers, J. S. Kissel, M. J. Szczepańczyk, and S. Vitale, *Phys. Rev. D* **101**, 042003 (2020), arXiv:1911.09083 [gr-qc].

[21] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Roheni, S. Allen, M. Cabero, K. Crowston, A. K. Katsagellos, S. L. Larson, T. K. Lee, C. Lintott, T. B. Littenberg, A. Lundgren, C. Østerlund, J. R. Smith, L. Trouille, and V. Kalogera, *Classical and Quantum Gravity* **34**, 064003 (2017), arXiv:1611.04596 [gr-qc].

[22] J. Merritt, B. Farr, R. Hur, B. Edelman, and Z. Doctor, arXiv e-prints , arXiv:2108.12044 (2021), arXiv:2108.12044 [gr-qc].

[23] D. e. a. Davis, *Classical and Quantum Gravity* **38**, 135014 (2021), arXiv:2101.11673 [astro-ph.IM].

[24] 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, *Classical and Quantum Gravity* **37**, 175001 (2020), arXiv:2001.01462 [gr-qc].

[25] D. Davis, L. V. White, and P. R. Saulson, *Classical and Quantum Gravity* **37**, 145001 (2020), arXiv:2002.09429 [gr-qc].

[26] R. Essick, P. Godwin, C. Hanna, L. Blackburn, and E. Katsavounidis, arXiv e-prints , arXiv:2005.12761 (2020), arXiv:2005.12761 [astro-ph.IM].

[27] E. e. a. Cuoco, arXiv e-prints , arXiv:2005.03745 (2020), arXiv:2005.03745 [astro-ph.HE].

[28] M. Razzano and E. Cuoco, *Classical and Quantum Gravity* **35**, 095016 (2018), arXiv:1803.09933 [gr-qc].

[29] N. Mukund, S. Abraham, S. Kandhasamy, S. Mitra, and N. S. Philip, *Physical Review D* **95**, 104059 (2017), arXiv:1609.07259 [astro-ph.IM].

[30] G. A. Valdes Sanchez, *Data Analysis Techniques for Ligo Detector Characterization*, Ph.D. thesis, The University of Texas at San Antonio (2017).

[31] T. J. Massinger, *Detector characterization for advanced LIGO*, Ph.D. thesis, Syracuse University (2016).

[32] L. K. Nuttall, T. J. Massinger, J. Areeda, J. Betzwieser, S. Dwyer, A. Effler, R. P. Fisher, P. Fritschel, J. S. Kissel, A. P. Lundgren, D. M. Macleod, D. Martynov, J. McIver, A. Mullavey, D. Sigg, J. R. Smith, G. Vajente, A. R. Williamson, and C. C. Wipf, *Classical and Quantum Gravity* **32**, 245005 (2015), arXiv:1508.07316 [gr-qc].

[33] 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, *Physical Review D* **88**, 062003 (2013), arXiv:1303.6984 [astro-ph.IM].

[34] D. MacLeod, Improving the sensitivity of searches for gravitational waves from compact binary coalescences, Ph.D. thesis, Cardiff University (United Kingdom) (2013).

[35] J. e. a. Aasi, *Classical and Quantum Gravity* **29**, 155002 (2012), arXiv:1203.5613 [gr-qc].

[36] N. Christensen, LIGO Scientific Collaboration, and Virgo Collaboration, *Classical and Quantum Gravity* **27**, 194010 (2010).

[37] T. Isogai, L. S. Collaboration, V. Collaboration, et al., in *Journal of Physics Conference Series*, Vol. 243 (IOP Publishing, 2010) p. 012005.

[38] S. Mukherjee, R. Obaid, and B. Matkarimov, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 243 (2010) p. 012006.

[39] L. e. a. Blackburn, *Classical and Quantum Gravity* **25**, 184004 (2008), arXiv:0804.0800 [gr-qc].

[40] D. Sigg, R. Bork, and J. Zweizig, in *The Ninth Marcel Grossmann Meeting*, edited by V. G. Gurzadyan, R. T. Jantzen, and R. Ruffini (2002) pp. 1841–1842.

[41] R. Gurav, B. Barish, G. Vajente, and E. E. Papalexakis, in *AAI 2020 Fall Symposium on Physics-Guided AI to Accelerate Scientific Discovery*, M. Puma, H. Radkins, C. Ramet, S. Henning, and C. Ruet, p. 1 (2020).

[42] R. E. Colgan, K. R. Corley, Y. Lau, I. Bartos, J. N. Wright, Z. Márka, and S. Márka, **101**, 102003.

[43] 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,” (2022), arXiv:2202.13486 [cs.LG].

[44] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, M. Adamo, C. Adams, T. Adams, P. Addesso, et al., *Classical and Quantum Gravity* **33**, 134001 (2016), arXiv:1602.03844 [gr-qc].

[45] D. Davis and M. Walker, *Galaxies* **10** (2022), 10.3390/galaxies10010012.

[46] J. D. Merritt, B. Farr, R. Hur, B. Edelman, and Z. Doctor, *Phys. Rev. D* **104**, 102004 (2021), arXiv:2108.12044 [gr-qc].

[47] P. e. a. Nguyen, *Classical and Quantum Gravity* **38**, 145001 (2021), arXiv:2101.09935 [astro-ph.IM].

[48] L. White and LIGO Team, in *APS April Meeting Abstracts*, APS Meeting Abstracts, Vol. 2019 (2019) p. K01.063.

[49] D. Davis, Improving the sensitivity of Advanced LIGO through data-driven methods, Ph.D. thesis, Syracuse University, United States (2019).

[50] F. Robinet, “Omicron: An Algorithm to Detect and Characterize Transient Noise in Gravitational-Wave Detectors,” <https://tts.ego-gw.it/ql/?c=10651> (2015).

[51] H. Zou and T. Hastie, *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**, 301 (2005).

[52] J. Rollins, Multimessenger Astronomy with Low-Latency Detectors, Ph.D. thesis, Columbia University (2011).

[53] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis, *Classical and Quantum Gravity* **21**, S1809 (2004), arXiv:gr-qc/0412119 [gr-qc].

[54] S. Soni and T. L. S. Collaboration, **38**, 025016.

[55] S. Bianchi, A. Longo, G. Valdes, G. González, and W. Plastino, arXiv e-prints , arXiv:2107.07565 (2021), arXiv:2107.07565 [astro-ph.IM].

[56] C. D. Austin, Measurements and mitigation of scattered light noise in LIGO, Ph.D. thesis, Louisiana State University, United States (2020).

[57] J. Harms, J. Lee, S. Sankar, P. R. Saulson, and J. R. Smith, *Journal of the Optical Society of America A* **29**, 1722 (2012), arXiv:1204.2528 [physics.optics].

[58] The Virgo Collaboration, arXiv e-prints , arXiv:1108.1598 (2011), arXiv:1108.1598 [gr-qc].

[59] J.-Y. Vinet, V. Brisson, S. Braccini, I. Ferrante, L. Pinard, F. Bondu, and E. Tournié, *Phys. Rev. D* **56**, 6085 (1997).

[60] P. Saha, Noise analysis of a suspended high power Michelson interferometer, Ph.D. thesis, MASSACHUSETTS INSTITUTE OF TECHNOLOGY (1997).

[61] Y. Asali and Z. Marka, “Effect of commissioning break on scattering glitches, ligo document t2000052,”.

[62] F. Maticichard, B. Lantz, R. Mittleman, K. Mason, J. Kissel, B. Abbott, S. Biscans, J. McIver, R. Abbott, S. Abbott, E. Allwine, S. Barnum, J. Birch, C. Celerier, D. Clark, D. Coyne, D. DeBra, R. DeRosa, M. Evans, S. Foley, P. Fritschel, J. A. Giaime, C. Gray, G. Grabeel, J. Hanson, C. Hardham, M. Hillard, W. Hua, C. Kucharczyk, M. Landry, A. Le Roux, V. Lhuillier, D. Macleod, M. Macinnis, R. Mitchell, B. O'Reilly, D. Ottaway, S. Punturo, M. Puma, H. Radkins, C. Ramet, M. Robinson, L. Ruet, P. Sarin, D. Shoemaker, A. Stein, J. Thomas, M. Vargas, K. Venkateswara, J. Warner, and S. Wen, *Classical and Quantum Gravity* **32**, 185003 (2015), arXiv:1502.06300 [physics.ins-det].

[63] D. Sigg, in *Interferometry XVIII*, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9960, edited by K. Creath, J. Burke, and A. Albertazzi Gonçalves (2016) p. 996009.

[64] J. G. Rollins, *Review of Scientific Instruments* **87**, 094502 (2016).

[65] A. Buikema, C. Cahillane, G. Mansell, C. Blair, R. Abbott, C. Adams, R. Adhikari, A. Ananyeva, S. Appert, K. Arai, et al., **102**, 062003, 2008.01301.

[66] B. Shapiro, “Overview of advanced ligo suspensions, ligo document g1100866,”.

[67] B. N. Shapiro, R. Adhikari, J. Driggers, J. Kissel, B. Lantz, J. Rollins, and K. Youcef-Toumi, *Classical and Quantum Gravity* **32**, 015004 (2015).

[68] S. M. Aston, M. A. Barton, A. S. Bell, N. Beveridge, B. Bland, A. J. Brummitt, G. Cagnoli, C. A. Cantley, L. Carbone, A. V. Cumming, L. Cunningham, R. M. Cutler, D. Heffner, J. Heptonstall, G. D. Hammond, K. Haughian, T. M. Hayler, A. Heptonstall, J. Heefner, D. Hoyland, J. Hough, R. Jones, J. S. Kissel, R. Kumar, N. A. Lockerbie, D. Lodhia, I. W. Martin, P. G. Murray, J. O'Dell, M. V. Plissi, S. Reid, J. Romie, N. A. Robertson, S. Rowan, B. Shapiro, C. C. Speake, K. A. Strain, K. V. Tokmakov, C. Torrie, A. A. van Veggel, A. Vecchio, and T. Ward, *Classical and Quantum Gravity* **29**, 235004 (2012).

[69] 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, *Classical and Quantum Gravity* **29**, 115005 (2012), arXiv:1205.5643 [gr-qc].

[70] A. V. Cumming, A. S. Bell, L. Barsotti, M. A. Barton, G. Cagnoli, D. Cook, L. Cunningham, M. Evans, R. Jones, R. Kumar, R. Mittleman, N. A. Robertson, S. Rowan, B. Shapiro, K. A. Strain, K. Tokmakov, C. Torrie, and A. A. van Veggel, *Classical and Quantum Gravity* **29**, 235005 (2012), arXiv:1205.5644 [gr-qc].

tum Gravity **29**, 035003 (2012).

[71] N. A. Robertson and E. Majorana, in *Advanced Interferometric Gravitational-Wave Detectors*, Vol. **21**, E. Essam Schwartz, Gravitational Wave Detectors, Edited by Reitze D et al. (2019) pp. 423–457.

[72] L. Barsotti, M. Evans, and P. Fritschel, *Classical and Quantum Gravity* **27**, 084026 (2010).

[73] 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, 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 (2004) pp. 194–205.

[74] R. Abbott, R. Adhikari, G. Allen, D. Baglino, C. Campbell, D. Coyne, E. Daw, D. DeBra, J. Faludi, P. Fritschel, A. Ganguli, J. Giaime, M. Hammond, C. Hardham, G. Harry, W. Hua, L. Jones, J. Kern, B. Lantz, K. Lilienkamp, K. Mailand, K. Mason, R. Mittleman, S. Nayfeh, D. Ottaway, J. Phinney, W. Rankin, N. Robertson, R. Scheffler, D. H. Shoemaker, S. Wen, M. Zucker, and L. Zuo, *Classical and Quantum Gravity* **21**, S915 (2004).

[75] E. Essam Schwartz, *Gravitational Wave Detectors*, Edited by Reitze D et al. (2019) pp. 235007 (2020), arXiv:2007.12847 [physics.ins-det].

[76] M. e. a. Walker, *Review of Scientific Instruments* **88**, 124501 (2017), arXiv:1702.04701 [astro-ph.IM].

[77] H. P. Daveloza, M. Afrin Badhan, M. Diaz, K. Kawabe, P. N. Konverski, M. Landry, and R. L. Savage, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 363 (2012) p. 012007.

[78] K. A. Strain and B. N. Shapiro, *Review of Scientific Instruments* **83**, 044501-044501-9 (2012).

[79] 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, *Nuclear Instruments and Methods in Physics Research A* **598**, 737 (2009).

[80] B. Canuel, E. Genin, G. Vajente, and J. Marque, *Optics Express* **21**, 10546 (2013).

[81] D. J. Ottaway, P. Fritschel, and S. J. Waldman, *Optics Express* **20**, 8329 (2012).