

# Multivariate Time Series Clustering for Environmental State Characterization of Ground-Based Gravitational-Wave Detectors

Rutuja Gurav

*Computer Science & Engineering  
University of California, Riverside  
Riverside, CA, USA  
rutuja.gurav@email.ucr.edu*

Isaac Kelly

*Physics  
University of Dallas  
Dallas, TX, USA  
ikelly@udallas.edu*

Pooyan Goodarzi

*Physics & Astronomy  
University of California, Riverside  
Riverside, CA, USA  
pooyan.goodarzi@email.ucr.edu*

Anamaria Effler

*LIGO Laboratory  
LIGO Livingston Observatory  
Livingston, LA, USA  
aeffler@ligo-la.caltech.edu*

Barry Barish

*Physics & Astronomy  
University of California, Riverside  
Riverside, CA, USA  
barry.barish@ucr.edu*

Evangelos E. Papalexakis

*Computer Science & Engineering  
University of California, Riverside  
Riverside, CA, USA  
epapalex@cs.ucr.edu*

Jonathan W. Richardson

*Physics & Astronomy  
University of California, Riverside  
Riverside, CA, USA  
jonathan.richardson@ucr.edu*

**Abstract**—Gravitational-wave observatories like LIGO are large-scale, terrestrial instruments housed in infrastructure that spans a multi-kilometer geographic area and which must be actively controlled to maintain operational stability for long observation periods. Despite exquisite seismic isolation, they remain susceptible to seismic noise and other terrestrial disturbances that can couple undesirable vibrations into the instrumental infrastructure, potentially leading to control instabilities or noise artifacts in the detector output. It is, therefore, critical to characterize the seismic state of these observatories to identify a set of temporal patterns that can inform the detector operators in day-to-day monitoring and diagnostics. On a day-to-day basis, the operators monitor several seismically relevant data streams to diagnose operational instabilities and sources of noise using some simple empirically-determined thresholds. It can be untenable for a human operator to monitor multiple data streams in this manual fashion and thus a distillation of these data-streams into a more human-friendly format is sought. In this paper, we present an end-to-end machine learning pipeline for features-based multivariate time series clustering to achieve this goal and to provide actionable insights to the detector operators by correlating found clusters with events of interest in the detector.

## I. INTRODUCTION

In the last nine years, the Laser Interferometer Gravitational-wave Observatory (LIGO) [1] and the European Virgo observatory [4] have established gravitational waves as a new observational probe of the universe. Gravitational waves, disturbances in the geometry of spacetime generated by the acceleration of matter, are a longstanding prediction of general relativity. The first direct detection of gravitational waves from two merging black holes in 2015 by LIGO [7] opened a new window on the universe. A subsequent observation of a binary neutron star merger with gravitational waves [8] and every band of the electromagnetic spectrum [10], [11] launched a new *multi-messenger* era of astronomy. The current generation

of detectors have now observed a wide variety of merger events involving black holes and neutron stars [2], [9], [28], [30]. The insights provided by these gravitational-wave signals are being used to address long-standing questions in astrophysics and fundamental physics.

Gravitational-wave detectors like LIGO are large-scale, terrestrial instruments housed in an infrastructure which sprawls a large geographic area. The two LIGO detectors, located in Hanford, WA and Livingston, LA, each consist of a 4-km-long Michelson laser interferometer whose sensitivity is further enhanced using multiple internal laser cavities [16]. The length and angular degrees of the individual laser cavities, as well as the interferometer as a whole, must be sensed and actively controlled in order to maintain operational stability for long observation periods. Despite exquisite seismic isolation of the detectors' optics, they remain susceptible to seismic noise and other terrestrial disturbances that can couple undesirable vibrations into the instruments' infrastructure. Figure 1 shows an example of some of the environmental disturbances regularly recorded by sensors at the LIGO sites.

By introducing physical motions between the interferometers' optics, environmental disturbances limit LIGO's sensitivity to gravitational waves at the lower end of its sensitive band, below 20 Hz. Environmental noise also poses a serious challenge to the operational stability and data quality of the detectors. For example, environmental conditions may account for many detector *glitches*, nonstationary noise bursts of largely unknown origin, that severely degrade the detector sensitivity during their duration [12], [13], [17]. Glitches contaminate the astrophysical data streams, confusing the gravitational-wave search pipelines and hindering the timely issuance of real-time alerts for electromagnetic follow-up observations. Elevated environmental noise is also known to cause *lock*

losses, control failures occurring when a disturbance causes the laser cavities to become driven too far from their resonant operating points [35]. Because lock re-acquisition typically requires 30 minutes to complete, lock losses limit the duty cycle for multi-detector observations, which are essential for precise sky localization of gravitational-wave events.

As a result, LIGO is extremely interested in understanding the emergence and effects of external disturbances and monitoring their behavior over time. For example, periods of increased glitch rates have been anecdotally associated with elevations of certain types of environmental noise. In fact, internally, detector commissioners already have heuristic means for manually monitoring such behavior. The motivation of this work is a result of direct collaboration with LIGO in order to automate, and extend, this endeavor. Our objective is to distill the information from a large number of heterogeneous environmental sensors, distributed across the 4-km LIGO sites, into a single environment “state” word that can be continuously recorded and tracked over time. Beyond automation, this tracking has the potential to reveal new, previously unrecognized associations between specific environmental conditions and detector anomalies (e.g., periods of increased glitch rates or controls instabilities). Such associations can provide actionable insight into the physical nature of the anomalies, directly guiding detector commissioning.

In this paper, we present an end-to-end multivariate time series analysis pipeline built to characterize the environmental state of ground-based gravitational-wave observatories like LIGO. The task at hand is to *identify known seismic and other environmental phenomena (e.g., earthquakes, anthropogenic activity) in measurements made by the network of sensors deployed across the detector sites and assign an interpretable label to each point in time identifying combinations of phenomena and the location(s) where they are active*. The major aspects of this work include:

- **End-to-end data science pipeline:** A major contribution of this work is identifying what data sources to monitor and collect and how to translate the task, defined by a real LIGO need, to a data science pipeline. Given our active collaboration with LIGO, this pipeline has the potential to be deployed at the sites as a powerful diagnostic tool.
- **Dataset:** We accompany our work with a public release of a relevant dataset that can foster further research in this direction. The dataset includes all of the LIGO time series data used in the development of our pipeline, hosted and provided by the Gravitational Wave Open Science Center (GWOSC) [3], [29].<sup>1</sup>

The rest of this paper is organized as follows. In Section II we provide a brief background on state-of-the-art gravitational-wave detectors, followed by a review of the existing work related to environmental state characterization. Section III presents our proposed data science pipeline. We then present results on a real LIGO dataset in Section IV to demonstrate the

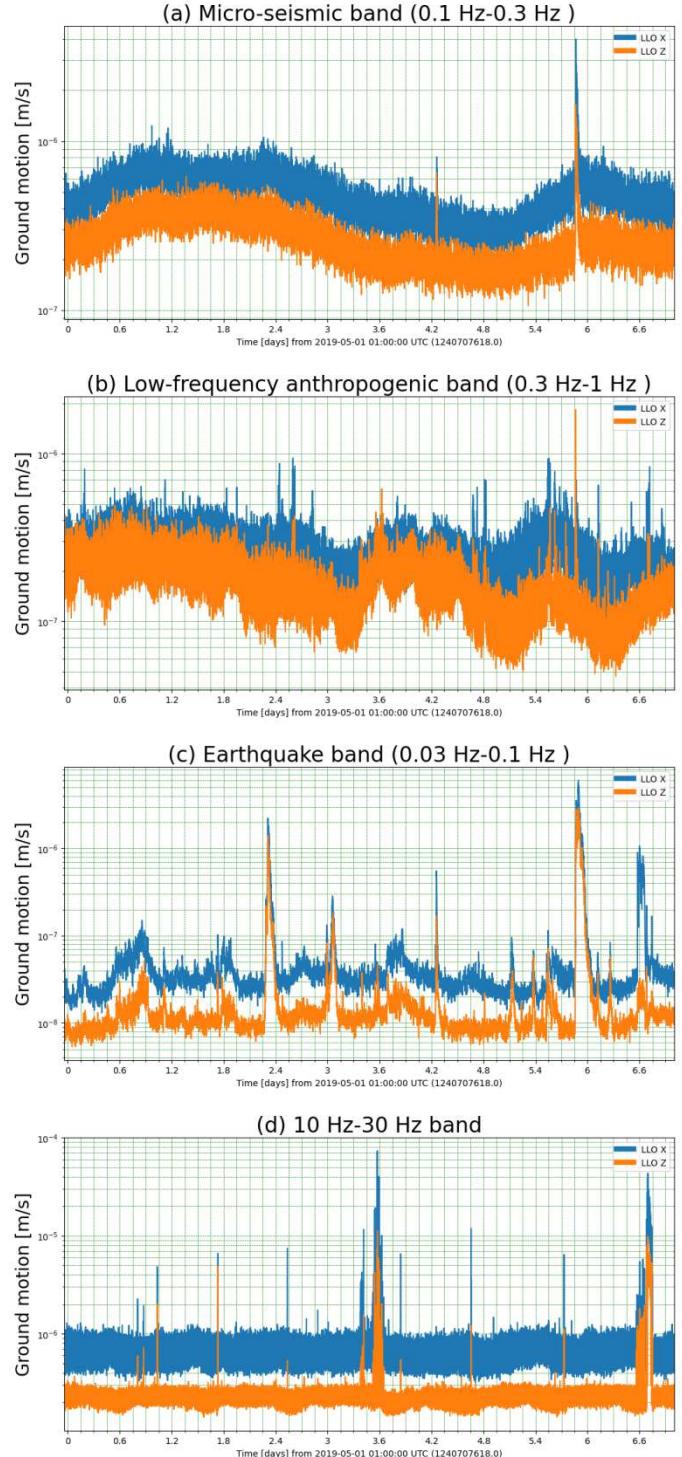


Fig. 1. A one-week band-limited sample of root-mean-square (RMS) ground motion data recorded by seismometers at the LIGO Livingston site. Each frequency band is associated with a set of different physical causes. (a) Micro-seismic frequency band (0.1-0.3 Hz) is mostly sensitive to ground motion caused by oceanic waves and has a characteristic time scale of multiple days [21]. (b) Low-frequency anthropogenic band (0.3-1 Hz) is correlated to various human related activities and the tides. (c) Earthquake band (0.03-0.1 Hz) captures ground motions mostly due to earthquakes and wind. (d) The 10-30 Hz frequency band is sensitive to ground motion due to mechanical vibrations of equipment at the LIGO sites, such as the HVAC system [31].

<sup>1</sup>Data available at: <https://gwosc.org/O3/trend/>

effectiveness of this pipeline. Finally, we conclude this paper in Section V.

## II. BACKGROUND AND MOTIVATION

### A. Gravitational-Wave Detectors

Gravitational-wave detectors are multi-kilometer-scale instruments ranking among the largest and most complex scientific facilities in the world. In addition to the main channel sensitive to spacetime strain, where gravitational-wave signals are observed, each LIGO detector has over 100,000 auxiliary channels which monitor the operation of each subsystem and the seismic, acoustic, and electromagnetic environment. LIGO’s physical environmental monitoring (PEM) system [5], [19], [31] consists of a distributed network of accelerometers, seismometers, microphones, magnetometers, power-mains voltage monitors, radio-frequency receivers, cosmic-ray detectors, and wind, temperature, and humidity sensors. This wealth of data presents a unique and largely untapped opportunity: *Can we leverage these vast amounts of data to improve our understanding of the relationship between changing environmental conditions and their impact on the detector’s performance?*

### B. Relation to Previous Work

In the context of environmental state characterization, a prior pipeline has been developed within LIGO to provide real-time seismic predictions to the interferometer operators. Seismon [15] is an earthquake early-warning system deployed at both LIGO detector sites. It uses near real-time earthquake alerts provided by the U.S. Geological Survey (USGS) and the National Oceanic and Atmospheric Administration (NOAA) to estimate the time of arrival and amplitude of the surface waves of earthquakes from around the globe at each detector site. Based on the predicted amplitude and direction of the surface waves, Seismon provides a machine-learning-based estimate of the probability that the incoming seismic event will cause a lock loss.

Seismon has enhanced the operational stability and up-time of the LIGO detectors by providing operators with the information necessary for on-site decision making to transition the detector to a more earthquake-robust operational mode. “Earthquake mode” is an alternative controls strategy which allows the detectors’ length servos to handle larger seismic disturbances, but at the expense of increased instrumental noise (reduced astrophysical sensitivity). The work presented here plays a highly complementary role to Seismon. Rather than making forward-looking predictions for real-time decision making, our pipeline is designed to characterize the *current* environmental state and can potentially include many other forms of disturbances, beyond seismic phenomena, as well.

Furthermore, in a broader context of state characterization, various detector characterization methods are used to generate data quality products from both the main strain channel and the auxiliary channels. These products take the form of bit-vector flags that indicate time segments of specific quality within each interferometer [17]. For instance, there are “observing

mode” flags at both sites generated by detector operators, indicating time segments in which the interferometer is locked and is reliable for astrophysical data inference. In this study, we only utilized the segments flagged in this manner, defining the detector’s status as being in observing mode.

### C. Related Time Series Work

There exist a number of relevant works in the time series mining and learning literature that have tackled problems that are similar to the one at hand. Matsubara *et al.* [27] has developed AutoPlait, a co-evolving time series mining tool that can identify a general set of patterns among a collection of time series that are related. There is even earlier work by Papadimitriou *et al.* [33] that attempts to do so, with the additional constraint that the data is seen as a stream, which introduces computational challenges.

Beyond the general-purpose identification of patterns in multiple time series, there have been various problem definitions and associated solutions which can fit our scenario and are also highly related to each other. These include regime shifts in multivariate time series [26], change detection [23], and multivariate time series segmentation [20]. Broadly, these techniques all seek to identify periods of correlated behavior across a collection of time series and points in time where those periods change from one category to another. Additionally, works that detect anomalies in multivariate time series [18], [24], [32], [37], while not directly addressing the problem definition at hand, computationally require a similar approach. Here, the interest is in identifying irregular patterns that far exceed the “normal” behavior, which the rest of the related works seek to characterize.

## III. PROPOSED PIPELINE

In this section, we describe our pipeline consisting of three modules:

- 1) Dataset Creation
- 2) Modeling
- 3) Downstream Evaluation

Our code-base is available at UC Riverside’s git repository.<sup>2</sup>

### A. Dataset Creation

*a) Data:* There are numerous seismometers deployed across the detector sites to monitor seismic activity. Figure 2 (top) shows a readout from one of these sensors. There are known seismic phenomena that manifest in certain frequency ranges, Figure 2 (bottom) shows three time series which are obtained by bandpass filtering the seismometer readout to isolate the following physically-motivated frequency ranges.

- 1) **0.03-0.1 Hz Earthquake band:** This frequency band is sensitive to ground motion due to earthquakes.
- 2) **0.1-0.3 Hz Microseism band:** This frequency band is sensitive to ground motion due to ocean waves beating against the shore, dominated by the Pacific Ocean for

<sup>2</sup>[https://git.ligo.org/uc\\_riverside/state-characterization](https://git.ligo.org/uc_riverside/state-characterization)

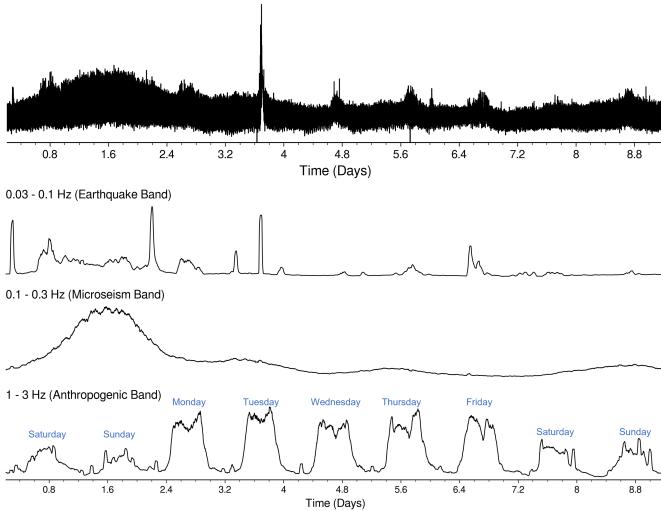


Fig. 2. Top: Example time series data from one of the numerous seismometers deployed across the LIGO sites. Bottom: The same signal bandpass-filtered to select three physically-motivated frequency bands corresponding to known seismic phenomenon.

the LIGO Hanford detector and the Atlantic Ocean and Gulf of Mexico for the LIGO Livingston detector.

3) **1-3 Hz Anthropogenic band:** This frequency band is sensitive to ground motion due to daily human activity, such as heavy traffic on nearby roads, passing trains, and logging or construction close to a LIGO site.

Although LIGO data channels have various sample rates depending on the physical quantity they measure and the specific sensors used for measurement, in this study, we used the band-limited RMS-averaged (BLRMS) second-trend data from each channel. Therefore, each data point is generated by calculating the root mean square of one-second segments of the recorded data.

In the current analysis, we used the data from a one month period during the O3b run. The utilized sensors are five tri-axial seismometers (STS) located at different positions of the LIGO Livingston's Internal Seismic Isolation subsystem (L1-ISI), namely at ETMX, ETMY, ITMX, ITMY, and HAM5 [1]. Each sensor has three orthogonal axes and is band-limited to six physically motivated frequency bands. The total data volume for the 90 second-trend L1-ISI-STS-BLRMS channels used in this analysis amounts to 20 GiB.

## B. Modeling

In our approach, given that (1) we require near real-time computations that can provide immediate insights to the LIGO operators and (2) we are expecting domain experts to deploy and fine-tune our tool, we would like to start from a simple approach with a minimal number of hyperparameters. As a result, even though (as we outlined in the related work above) there exist a number of off-the-shelf methods that could be adapted for the purposes of this work, we instead opt for a light-weight clustering-based approach which can run fast and requires the definition of only two hyperparameters:

the duration of window segments we are considering and the number of clusters that correspond to the sought-after states. The window segment size is a parameter that our domain expert collaborators are very confident setting up given their empirical knowledge stemming from working with the detector. The number of clusters can be relatively easily narrowed down by popular heuristics [36], minimizing the time spent in hyperparameter tuning by the operators.

Our proposed algorithm for identifying different environmental states of the detector based on the subset of the channels we have selected is as follows:

- 1) We are given a batch of  $C$  time-series, corresponding to the different channels of interest. The current chosen length is one hour.
- 2) For a given input window size  $w$ , create  $C$  segments per non-overlapping window of the entire length.
- 3) For each set of  $C$  time-series corresponding to the same window, derive a set of statistical features. In this way, a given window becomes a data point. We experimented with (a) simple statistics such as mean and standard deviation, (b) tsfresh [14], and (c) catch22 [25]. In our use cases we found (a) to be sufficient, but if and when we introduce more complex patterns, more sophisticated features may become necessary.
- 4) Run k-means clustering on all windows/data points. We utilized the k-means++ [6] algorithm to find the initial seeds for the k-means clustering. We determine the number of clusters by intersecting the results with known intrinsic cluster validation indices [36], see 5. Scikit-learn [34] library was used for all of the modeling done in this analysis.
- 5) As a final step, we need to provide a human-understandable set of labels for our results. Instead of using the raw names of the channels involved, which are not necessarily always intuitive, even to an experienced operator, we create three different “replicas” of the channels: (a) Anthropogenic, (b) Microseism, and (c) Earthquake. They correspond to known thresholds that LIGO operators currently manually use to identify one of those events. We then take the centroid of a given cluster and compare it against the known thresholds, and label the cluster accordingly.

Figure 3 graphically illustrates this workflow.

## C. Downstream Evaluation

Obtaining ground truth for our task is a rather ill-defined problem. For some of the discovered states, we can confirm the presence or absence of an earthquake by cross-referencing the discovered states with USGS data. However, less well-defined and properly monitored states, such as ones caused by anthropogenic factors, are very hard to validate. Thus, in lieu of ground truth, we investigate if and to what degree the discovered states correlate with instances where it has been documented that the detector was witnessing a noise glitch or experienced a loss of lock.

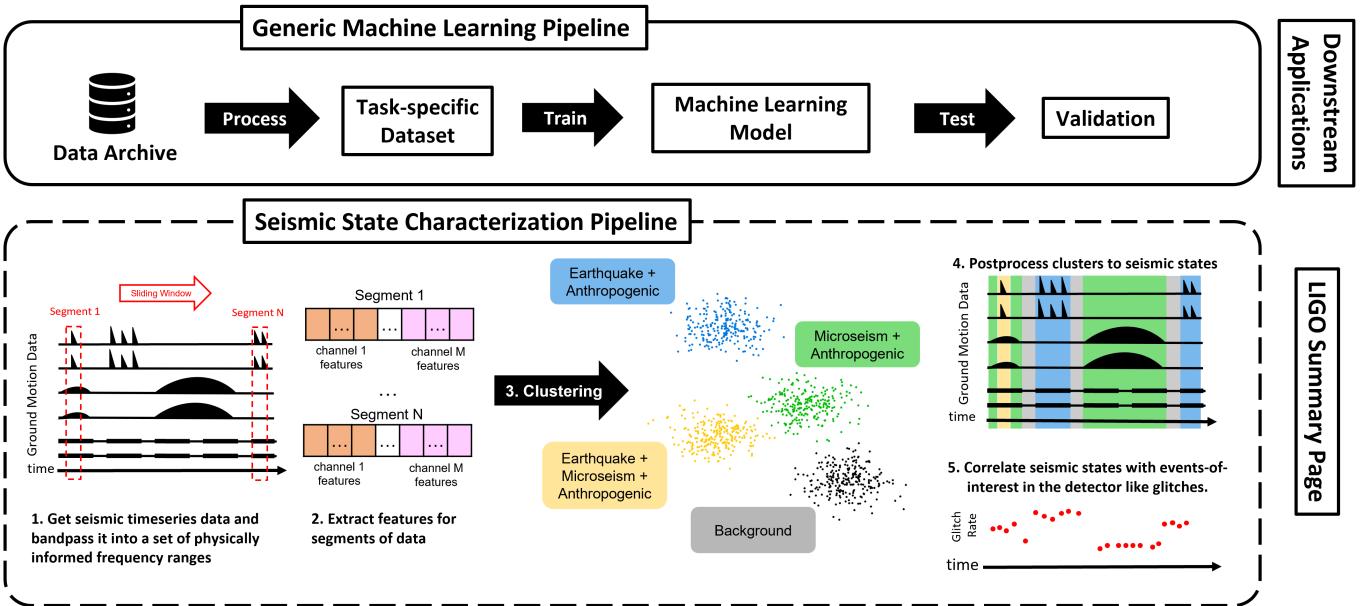


Fig. 3. Step-by-step workflow of our proposed pipeline. Each step is described in detail in the text of §III-B.

#### IV. RESULTS

In this section, we first demonstrate an example of running our tool end-to-end and what we envision the final deployed outcome would look like. Then, we present results that link certain discovered states with documented problematic detector states.

##### A. Indicative End-to-End Results

Figure 4 shows a snapshot of our results, starting from the different channel time series as they are band-passed in different frequency bands (top), to the raw clustering results (center), to the labeled states (bottom). We envision the bottom panel of the figure to be the final product of the tool shown to the operator in near real-time. Figure 5 shows an indicative set of results for identifying the number of clusters/states in a data sample. We find that different cluster validation indices have slight variation on the “best” number indicated. However, they seem to indicate a small range of admissible sets of states that are very feasible for the operator to iterate over and inspect.

##### B. Linking Discovered States to Glitches and Loss of Lock

One of the main goals of our proposed monitoring tool is to be able to diagnose problematic detector states, such as periods of controls instabilities and elevated noise glitch activity [22]. In order to do so, we conducted an experiment where we computed the expected glitch and loss of lock rates, assuming they occur randomly (i.e., have no physical relation to the identified environmental states), and then compare that expected rate to the observed rate per discovered state. The observed glitch rate was calculated using the publicly available Gravity Spy glitch classifications dataset [22], considering all the triggers with  $SNR > 7.5$  during the analysis period. We make an extremely fascinating observation: For some states,

the observed glitch rate far exceeds the expected rate, thus linking those states to those core detector issues. Figure 6 shows an example of such analysis, where some discovered states experience a much larger amount of glitches than would be randomly expected.

#### V. CONCLUSION & FUTURE WORK

In this paper, we present an instance of applied data science for detecting different environmental states of the LIGO detectors. This work has been directly motivated by working closely with LIGO commissioners and operators, understanding their needs, and translating them into a data science pipeline. We make interesting and important observations that link various discovered environmental states to documented issues faced by the detector, which is a positive step in the process of addressing those issues towards improving the detectors’ up-time and the quality of the science data.

In the near future, we will continue working closely with LIGO for subsequent deployment of our tool. Furthermore, we are interested in extending our framework in order to be able to accommodate *novel*, previously unrecognized environmental states. Finally, towards a highly-efficient deployment, we would like to explore the transition to a supervised model including known states discovered by our unsupervised tool and possibly obtain labeled data in a citizen science fashion, similar to the Gravity Spy project [22], [38] which does so for detecting glitches in the main channel of LIGO.

#### ACKNOWLEDGMENTS

This material is based upon work supported by NSF’s LIGO Laboratory which is a major facility fully funded by the National Science Foundation. The authors are grateful for computational resources provided by the LIGO Laboratory and

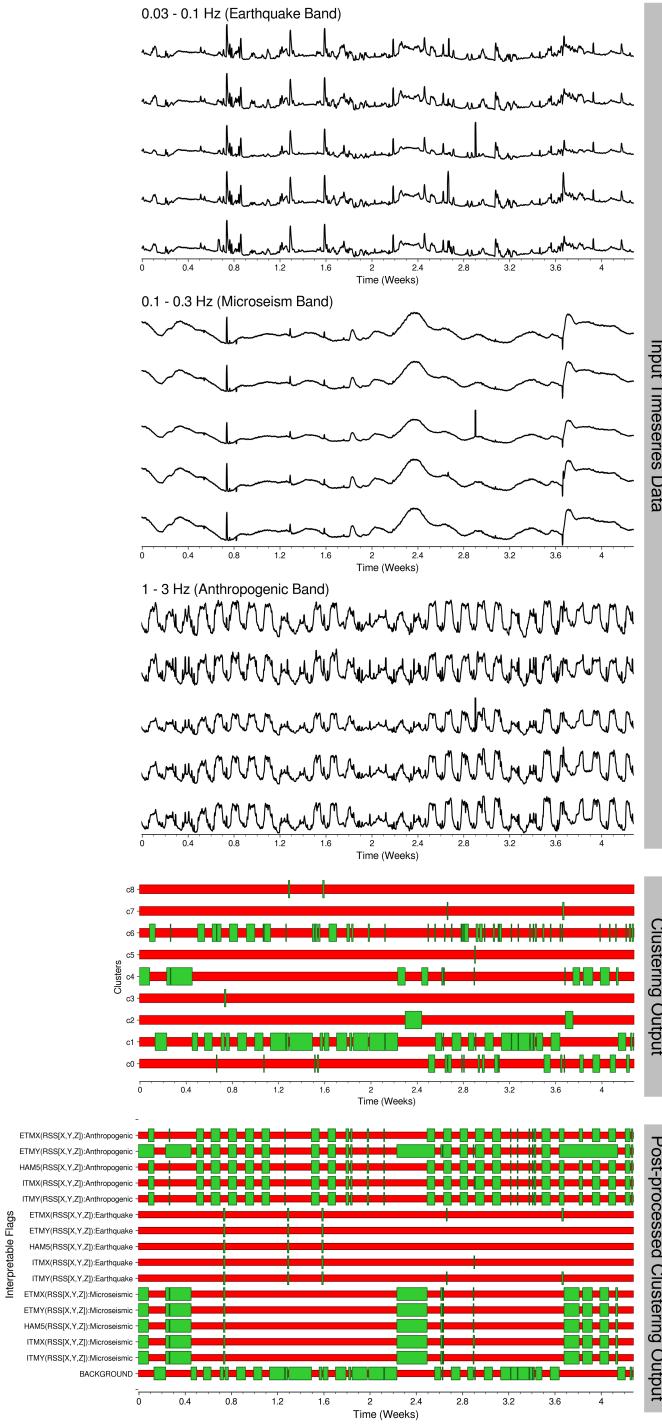


Fig. 4. Example of our end-to-end analysis. Top panel consists of a set of sample channels that we are running the model on. The middle panel is the output of the clustering model. Each line represents a different state and the green flags are the segments at which the detector is in that specific state. The bottom panel shows the states assigned by the field experts using a simple threshold. It can be seen that the model's output recovers the expert's expectations without supervision. For example, there are clusters that correspond to earthquakes, high microseism, and high anthropogenic noise.

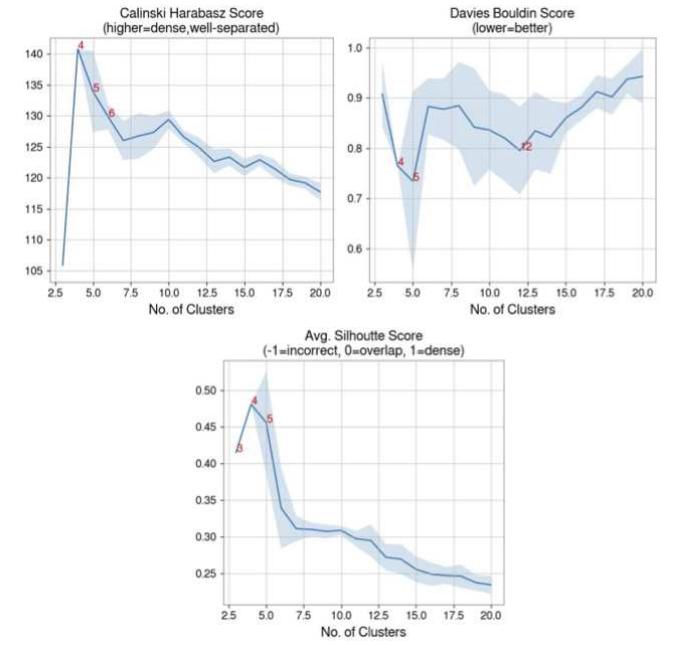


Fig. 5. Cluster validation indices [36] employed in order to identify a short range of admissible number of cluster values that the operator can iterate over. A grid search for [3-20] number of clusters was done and three standard clustering validation scores were calculated. The final number of clusters was determined using the correlation with the Glitch rates as an external validation metric.

supported by the National Science Foundation under Award Nos. PHY-0757058 and PHY-0823459. Research at UC Riverside was supported by the National Science Foundation under Award Nos. PHY-2141072 and IIS-2046086. This research has made use of data or software obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. This paper carries LIGO Document Number LIGO-P2400407.

## REFERENCES

- [1] J. Aasi et al. Advanced LIGO. *Class. Quant. Grav.*, 32:074001, 2015. doi: 10.1088/0264-9381/32/7/074001.
- [2] R. Abbott et al. GWTC-3: Compact Binary Coalescences Observed by LIGO and Virgo during the Second Part of the Third Observing Run. *Phys. Rev. X*, 13(4):041039, 2023. doi: 10.1103/PhysRevX.13.041039.
- [3] R. Abbott et al. Open Data from the Third Observing Run of LIGO, Virgo, KAGRA, and GEO. *Astrophys. J. Suppl.*, 267(2):29, 2023. doi: 10.3847/1538-4365/acdc9f.
- [4] F. Acernese et al. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Class. Quant. Grav.*, 32(2):024001, 2015. doi: 10.1088/0264-9381/32/2/024001.
- [5] Fausto Acernese, M Agathos, A Ain, S Albanesi, A Allocca, Alexandre Amato, T Andrade, N Andres, Marc Andrés-Carcasona, T Andrić, et al. The virgo o3 run and the impact of the environment. *Classical and quantum gravity*, 39(23):235009, 2022.
- [6] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. Technical Report 2006-13, Stanford InfoLab, June 2006.
- [7] B. P. Abbott et al. (LIGO Scientific Collaboration and Virgo Collaboration). Observation of Gravitational Waves from a Binary Black Hole Merger. *Physical Review Letters*, 116:061102, February 2016. doi: 10.1103/PhysRevLett.116.061102.

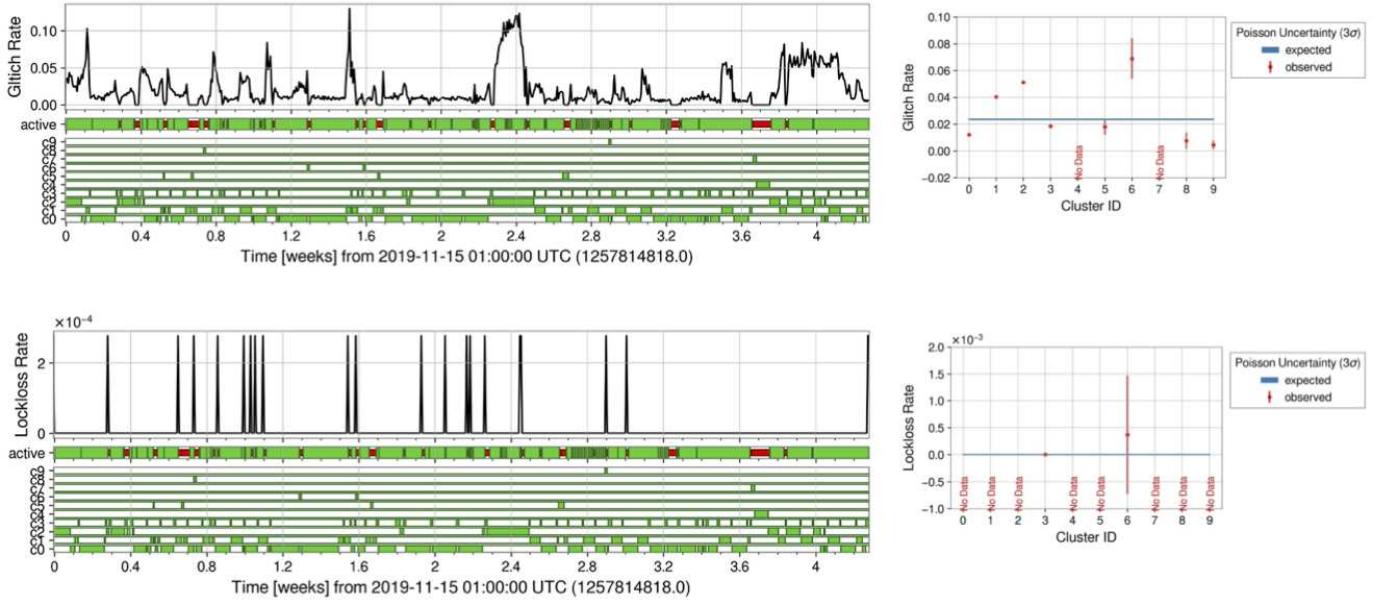


Fig. 6. Some of our discovered states experience a far higher number of glitches than expected, assuming random occurrence (top). This indicates that there is a physical connection between the identified environmental states and those core detector problems, so our pipeline can provide valuable diagnostic information to guide future detector improvements and commissioning.

- [8] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration). GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physical Review Letters*, 119:161101, October 2017. doi: 10.1103/PhysRevLett.119.161101.
- [9] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, and KAGRA Collaboration). Observation of Gravitational Waves from Two Neutron Star–Black Hole Coalescences. *The Astrophysical Journal Letters*, 915(1):L5, June 2021. doi: 10.3847/2041-8213/ac082e.
- [10] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, Fermi Gamma-ray Burst Monitor, and INTEGRAL). Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *The Astrophysical Journal Letters*, 848(2):L13, October 2017. doi: 10.3847/2041-8213/aa920c.
- [11] B. P. Abbott *et al.* (LIGO Scientific Collaboration, Virgo Collaboration, Fermi Gamma-ray Burst Monitor, INTEGRAL, IceCube, AstroSat Cadmium Zinc Telluride Imager Team, IPN, Insight-Hxmt, ANTARES, Swift, AGILE Team, 1M2H Team, Dark Energy Camera GW-EM, DES, DLT40, GRAWITA, Fermi-LAT, ATCA, ASKAP, Las Cumbres Observatory Group, OzGrav, DWF (Deeper Wider Faster Program), AST3, CAASTRO, VINROUGE, MASTER, J-GEM, GROWTH, JAGWAR, CaltechNRAO, TTU-NRAO, NuSTAR, Pan-STARRS, MAXI Team, TZAC Consortium, KU, Nordic Optical Telescope, ePESSTO, GROND, Texas Tech University, SALT Group, TOROS, BOOTES, MWA, CALET, IKI-GW Follow-up, H.E.S.S., LOFAR, LWA, HAWC, Pierre Auger, ALMA, Euro VLBI Team, Pi of Sky, Chandra Team at McGill University, DFN, ATLAS Telescopes, High Time Resolution Universe Survey, RIMAS, RATIR, SKA South Africa/MeerKAT). Multi-messenger observations of a binary neutron star merger. *The Astrophysical Journal Letters*, 848(2):L12, October 2017. doi: 10.3847/2041-8213/aa91c9.
- [12] L Blackburn, L Cadonati, S Caride, S Caudill, S Chatterji, N Christensen, J Dalrymple, S Desai, A Di Credico, G Ely, J Garofoli, L Goggin, G González, R Gouaty, C Gray, A Gretarsson, D Hoak, T Isogai, E Katsavounidis, J Kissel, S Klimenko, R A Mercer, S Mohapatra, S Mukherjee, F Raab, K Riles, P Saulson, R Schofield, P Shawhan, J Slutsky, J R Smith, R Stone, C Vorick, M Zanolin, N Zotov, and J Zweizig. The LSC glitch group: monitoring noise transients during the fifth LIGO science run. *Classical and Quantum Gravity*, 25(18):184004, September 2008. doi: 10.1088/0264-9381/25/18/184004.
- [13] M Cabero, A Lundgren, A H Nitz, T Dent, D Barker, E Goetz, J S Kissel, L K Nuttall, P Schale, R Schofield, and D Davis. Blip glitches in Advanced LIGO data. *Classical and Quantum Gravity*, 36(15):155010, July 2019. doi: 10.1088/1361-6382/ab2e14.
- [14] Maximilian Christ, Nils Braun, Julius Neuffer, and Andreas W Kempa-Liehr. Time series feature extraction on basis of scalable hypothesis tests (tsfresh—a python package). *Neurocomputing*, 307:72–77, 2018. doi: 10.1016/j.neucom.2018.03.067.
- [15] Michael Coughlin, Paul Earle, Jan Harms, Sébastien Biscans, Christopher Buchanan, Eric Coughlin, Fred Donovan, Jeremy Fee, Hunter Gabard, Michelle Guy, Nikhil Mukund, and Matthew Perry. Limiting the effects of earthquakes on gravitational-wave interferometers. *Classical and Quantum Gravity*, 34(4):044004, February 2017. doi: 10.1088/1361-6382/aa5a60.
- [16] D. V. Martynov *et al.* Sensitivity of the Advanced LIGO detectors at the beginning of gravitational wave astronomy. *Physical Review D*, 93:112004, June 2016. doi: 10.1103/PhysRevD.93.112004.
- [17] Derek Davis *et al.* LIGO detector characterization in the second and third observing runs. *Class. Quant. Grav.*, 38(13):135014, 2021. doi: 10.1088/1361-6382/abfd85.
- [18] Ailin Deng and Bryan Hooi. Graph neural network-based anomaly detection in multivariate time series. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 4027–4035, 2021. doi: 10.1609/aaai.v35i1.16523.
- [19] A Effler, R M S Schofield, V V Frolov, G González, K Kawabe, J R Smith, J Birch, and R McCarthy. Environmental influences on the LIGO gravitational wave detectors during the 6th science run. *Classical and Quantum Gravity*, 32(3):035017, January 2015. doi: 10.1088/0264-9381/32/3/035017.
- [20] Shaghayegh Gharabibi, Chin-Chia Michael Yeh, Yifei Ding, Wei Ding, Paul Hibbing, Samuel LaMunion, Andrew Kaplan, Scott E Crouter, and Eamonn Keogh. Domain agnostic online semantic segmentation for multi-dimensional time series. *Data mining and knowledge discovery*, 33:96–130, 2019. doi: 10.1007/s10618-018-0589-3.
- [21] J. A. Giaime, E. J. Daw, M. Weitz, R. Adhikari, P. Fritschel, R. Abbott, R. Bork, and J. Heefner. Feedforward reduction of the microseism disturbance in a long-base-line interferometric gravitational-wave detector. *Review of Scientific Instruments*, 74(1):218–224, January 2003. doi: 10.1063/1.1524717.
- [22] J Glanzer, S Banagiri, S B Coughlin, S Soni, M Zevin, C P L Berry, O Patane, S Bahaadini, N Rohani, K Crowston, V Kalogera, C Østerlund, L Trouille, and A Katsaggelos. Data quality up to the third observing run of Advanced LIGO: Gravity Spy glitch classifications. *Classical and*

*Quantum Gravity*, 40(6):065004, February 2023. doi: 10.1088/1361-6382/acb633.

[23] Bryan Hooi and Christos Faloutsos. Branch and border: Partition-based change detection in multivariate time series. In *Proceedings of the 2019 SIAM International Conference on Data Mining*, pages 504–512. SIAM, 2019. doi: 10.1137/1.9781611975673.57.

[24] Paloma Laguarta, Robin van der Laag, Melissa Lopez, Tom Dooney, Andrew L Miller, Stefano Schmidt, Marco Cavaglia, Sarah Caudill, Kurt Driessens, Joël Karel, et al. Detection of anomalies amongst ligo’s glitch populations with autoencoders. *Classical and Quantum Gravity*, 41(5):055004, 2024.

[25] Carl H Lubba, Sarab S Sethi, Philip Knaute, Simon R Schultz, Ben D Fulcher, and Nick S Jones. catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis. *Data Mining and Knowledge Discovery*, 33(6):1821–1852, 2019. doi: 10.1007/s10618-019-00647-x.

[26] Yasuko Matsubara and Yasushi Sakurai. Regime shifts in streams: Real-time forecasting of co-evolving time sequences. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1045–1054, 2016. doi: 10.1145/2939672.2939755.

[27] Yasuko Matsubara, Yasushi Sakurai, and Christos Faloutsos. Autoplait: Automatic mining of co-evolving time sequences. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 193–204, 2014. doi: 10.1145/2588555.2588556.

[28] Alexander H. Nitz, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Shichao Wu, Marlin Schäfer, Rahul Dhurkunde, and Collin D. Capano. 4-OGC: Catalog of Gravitational Waves from Compact Binary Mergers. *Astrophys. J.*, 946(2):59, 2023. doi: 10.3847/1538-4357/aca591.

[29] Gravitational Wave Open Science Center, O3 Second-trend Data from Seismometers, Wind Speed Monitors, and Accelerometers, Dec. 2024, doi: 10.7935/9yc2-5d96.

[30] Seth Olsen, Tejaswi Venumadhav, Jonathan Mushkin, Javier Roulet, Barak Zackay, and Matias Zaldarriaga. New binary black hole mergers in the LIGO-Virgo O3a data. *Phys. Rev. D*, 106(4):043009, 2022. doi: 10.1103/PhysRevD.106.043009.

[31] P Nguyen *et al.* Environmental noise in advanced LIGO detectors. *Classical and Quantum Gravity*, 38(14):145001, June 2021. doi: 10.1088/1361-6382/ac011a.

[32] Matteo Palenghi. *Time series anomaly detection for cern large-scale computing infrastructure*. PhD thesis, Milan, Polytech., 2020.

[33] Spiros Papadimitriou, Jimeng Sun, and Christos Faloutsos. Streaming pattern discovery in multiple time-series. 2005.

[34] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011.

[35] Jameson Rollins. Machine learning for lock loss analysis. LIGO Technical Report LIGO-G1701409-v1, July 2017.

[36] Erich Schubert. Stop using the elbow criterion for k-means and how to choose the number of clusters instead. *ACM SIGKDD Explorations Newsletter*, 25(1):36–42, 2023. doi: 10.1145/3606274.3606278.

[37] Sadaf Tafazoli and Eamonn Keogh. Matrix profile xxviii: Discovering multi-dimensional time series anomalies with k of n anomaly detection. In *Proceedings of the 2023 SIAM International Conference on Data Mining (SDM)*, pages 685–693. SIAM, 2023. doi: 10.1137/1.9781611977653.ch77.

[38] Michael Zevin, Corey B Jackson, Zoheyr Doctor, Yunan Wu, Carsten Österlund, L Clifton Johnson, Christopher PL Berry, Kevin Crowston, Scott B Coughlin, Vicky Kalogera, et al. Gravity spy: lessons learned and a path forward. *The European Physical Journal Plus*, 139(1):100, 2024. doi: 10.1140/epjp/s13360-023-04795-4.