

---

# Glitch Classification and Clustering for LIGO with Deep Transfer Learning

---

**Daniel George** \*

NCSA and Department of Astronomy  
University of Illinois at Urbana-Champaign  
dgeorge5@illinois.edu

**Hongyu Shen** \*

NCSA and Department of Statistics  
University of Illinois at Urbana-Champaign  
hongyu2@illinois.edu

**E. A. Huerta**

NCSA  
University of Illinois at Urbana-Champaign  
eliu@illinois.edu

## Abstract

The detection of gravitational waves with LIGO and Virgo requires a detailed understanding of the response of these instruments in the presence of environmental and instrumental noise. Of particular interest is the study of anomalous non-Gaussian noise transients known as glitches, since their high occurrence rate in LIGO/Virgo data can obscure or even mimic true gravitational wave signals. Therefore, successfully identifying and excising glitches is of utmost importance to detect and characterize gravitational waves. In this article, we present the first application of *Deep Learning* combined with *Transfer Learning* for glitch classification, using real data from LIGO's first discovery campaign labeled by Gravity Spy, showing that knowledge from pre-trained models for real-world object recognition can be transferred for classifying spectrograms of glitches. We demonstrate that this method enables the optimal use of very deep convolutional neural networks for glitch classification given small unbalanced training datasets, significantly reduces the training time, and achieves state-of-the-art accuracy above 98.8%. Once trained via transfer learning, we show that the networks can be truncated and used as feature extractors for unsupervised clustering to automatically group together new classes of glitches and anomalies. This novel capability is of critical importance to identify and remove new types of glitches which will occur as the LIGO/Virgo detectors gradually attain design sensitivity.

## 1 Introduction

The LIGO [1] and Virgo [2] detectors are the largest and most sensitive interferometric detectors ever built. They can even sense changes thousands of times smaller than the diameter of a proton [1, 3]. These instruments have already detected multiple gravitational wave (GW) signals produced from coalescence of black holes [4, 5, 6, 7] as well as neutron star mergers [8, 9]. As LIGO/Virgo gradually attain design sensitivity, they will transition into an astronomical observatory that will routinely detect new GW sources, providing insights into astrophysical events that cannot be seen through any other means [10, 11].

For LIGO/Virgo to realize their full potential, it is necessary to ensure that their sensing capabilities are not hindered by unwanted non-Gaussian noise transients, known as glitches, which contaminate

---

\*Equal contribution

GW data. There are extensive ongoing efforts on separating glitches from signals, and/or classifying them based on their characteristics, which is a non-trivial task requiring “intelligent” algorithms given that the glitches vary widely in duration, frequency range and morphology, spanning a wide distribution that is challenging to model accurately [12, 13, 14, 15, 16, 17, 18]. Furthermore, since the LIGO/Virgo detectors are undergoing commissioning between each observing run, we expect that new types of glitches will be identified as they attain reach sensitivity [3, 19, 20].

Accurately classifying glitches is essential for several reasons [14, 21]: a) This will prevent false GW detections due to coincident glitches across multiple detectors that closely mimic signals. b) Rapidly identifying and excising glitches will enhance detector sensitivity, and improve the significance of GW signals that are contaminated by glitches [8]. c) The LIGO/Virgo detectors have thousands of instrumental and environmental channels to monitor changes that occur due to environmental or hardware issues. By carefully tracking down these glitches, we aim to identify their source and eliminate them promptly to ensure that the data stream is usable for GW data analysis.

The complex and time-evolving nature of glitches makes them an ideal case study to apply machine learning algorithms. Deep learning algorithms have been recently applied for GW signal detection and parameter estimation [22] as well as for denoising LIGO data [23]. In this article, we focus on deep learning with Convolutional Neural Networks (CNNs) [24] for glitch classification, using spectrogram images computed from the time-series data as inputs. Recent efforts on this front include Gravity Spy, an innovative interdisciplinary project that provides an infrastructure for citizen scientists to label datasets of glitches from LIGO via crowd-sourcing [21]. Supervised classification algorithms based on this dataset were presented in the first Gravity Spy article [21] and were further discussed in [25]. These algorithms employed deep learning CNN models which were 4 layers deep, and achieved overall accuracies close to 97% for glitch classification. It was found, however, that glitch classes with very few labeled samples could not be classified with the same level of accuracy.

Here, we present *Deep Transfer Learning* as a new method for glitch classification that leverages pre-trained state-of-the-art CNNs used for object recognition, and fine-tunes them throughout all layers to accurately classify glitches after re-training on a small dataset of LIGO spectrograms. We show that this technique achieves state-of-the-art results for glitch classification with the Gravity Spy dataset, attaining above 98.8% overall accuracy and perfect precision-recall on 8 out of 22 classes, while significantly reducing the training time to a few minutes. Our results indicate that new types of glitches can be classified accurately given very few labeled examples with this technique. We also demonstrate that features learned from real-world images by very deep CNNs are directly transferable for the classification of spectrograms of time-series data from GW detectors, and possibly also for spectrograms in general, even though the two datasets are very dissimilar. The CNNs we use were originally designed for over 1000 classes of objects in ImageNet. Therefore, our algorithm can be easily extended to classify hundreds of new classes of glitches in the future, especially since this transfer learning approach requires only a few labeled examples of a new class. Furthermore, we outline how new classes of glitches can be automatically grouped together by using our trained CNNs as feature extractors for unsupervised or semi-supervised clustering algorithms.

## 2 Methods

The Gravity Spy crowd-sourcing project mobilizes citizen scientists to hand-label spectrograms obtained from LIGO time-series data after being shown only a few examples, which indicates that generic pattern recognizers developed in humans for real-world object recognition are also useful when distinguishing spectrograms of glitches. This motivated us to apply a similar approach referred as “transfer learning” in the machine learning literature.

Transfer learning is an essential ingredient for true artificial intelligence, where knowledge learned in one domain for some task (typically where there is a large amount of labeled data) can be transferred to another domain for a different task [26] where there may only be limited number of labeled examples. In the context of deep learning for classifying images, transfer learning can be performed by pre-training a deep CNN on a large and diverse dataset with well-established labels followed by modifying the final layer based on the number of required classes and then fine-tuning the weights on a different dataset of interest. It is well known that the initial layers of a CNN always learn to extract simple generic features (e.g., edges, corners, etc.), which are applicable to all types of images, whereas the final layers represent highly abstract and data-specific features [27]. Therefore, transfer

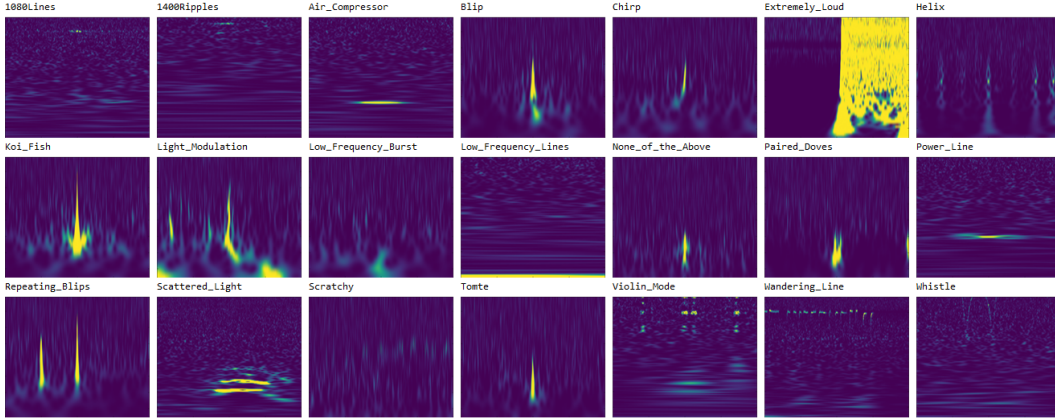


Figure 1: Classes of glitches in the Gravity Spy dataset from the first observing run of LIGO. The No\_Glitch class (periods of usual LIGO noise without transient glitches) is omitted in the figure. True GW signals from compact binary coalescence fall in the Chirp class. For each input, spectrograms with durations of 0.5s, 1.0s, 2.0s, and 4.0s are available. The objective is to predict the classes given the images.

learning is expected to result in higher accuracy and a faster training process compared to training the same CNNs from scratch, due to the shared features present in the initial layers.

To demonstrate the power of transfer learning for classifying glitches, we compare the performance of the most popular CNN models for object recognition, namely Inception [28] version 2 and 3, ResNet [29], and VGG [30], all of which were leading entries in recent ILSVRC competitions. These CNNs were pre-trained on a large dataset of images — i.e., ImageNet [31], which contains 1.2 million labeled images of real-world objects belonging to 1000 categories — over the course of 2 to 3 weeks using multiple GPUs by other research groups. We obtained the open-source weights from these models, and used them to initialize the CNNs, before fine-tuning (re-training) each model on our training dataset of glitches.

The Gravity Spy dataset, from the first observing run of LIGO, contains labeled spectrogram samples from 22 classes of glitches shown in Figure 1. We randomly split this dataset, containing about 8500 elements, into two parts such that approximately 80% of samples in each class was in the training set, and 20% of each class was in the testing set. These images were hand-labeled by citizen scientists participating in the Gravity Spy project, and the accuracy of the labeling was greatly enhanced by cross-validation techniques within the Gravity Spy infrastructure, also involving experts from the LIGO Detector Characterization team [21].

The final fully-connected layer in each CNN model was replaced with another fully-connected layer having 22 neurons corresponding to each glitch class. The softmax function is used as the final layer in each model to provide probabilities of each class as the outputs. We fine-tuned across all the layers since the dataset of glitches is very different from the objects in the ImageNet data.

### 3 Results

Both InceptionV2 and InceptionV3 achieved over 98% accuracy in fewer than 10 epochs of training (less than 20 minutes), VGG16 and VGG19 achieved over 98% accuracy within 30 epochs of training. In Table 1, we compare the results of these CNNs trained with the transfer learning method and that of the CNN models described in [21, 25] which were trained from scratch on the same training set for sufficient number of epochs. We found that our models consistently achieved over 98% accuracy for many epochs, thus indicating that the performance is robust, regardless of the stopping criteria, and therefore the model is not overfitting on the test set. Note that our models consistently under-performed with less than 98% accuracy when trained without transfer learning.

With InceptionV3, we achieved *perfect* precision and recall on 8 classes: 1080Lines, 1400Ripples, Air\_Compressor, Chirp, Helix, Paired\_Doves, Power\_Line, and Scratchy. With ResNet50,

Table 1: Accuracy of Classification on the Test Set

Neural Network	Top-1	Top-2	Top-3	Top-4	Top-5
Tuned-InceptionV3	98.84%	99.71%	99.88%	99.94%	100.00%
Tuned-InceptionV2	98.78%	99.59%	99.71%	99.94%	100.00%
Tuned-ResNet50	98.84%	99.71%	99.83%	99.94%	100.00%
Tuned-VGG16	98.15%	99.36%	99.71%	99.83%	99.88%
Tuned-VGG19	98.21%	99.31%	99.60%	99.71%	99.71%
CNN in [21, 25]	96.70%	98.32%	99.13%	99.31%	99.36%

This table lists the top-1 to top-5 accuracies for different CNNs on the testing set. We re-trained the 4 layer merged-view CNN model described in the publications [21, 25] from scratch on our same train-test dataset for a fair comparison. Each model was trained with sufficient numbers of iterations until the error on the validation set started increasing. Note that our Inception and ResNet models are capable of narrowing down any input to within 5 classes with 100.00% accuracy.

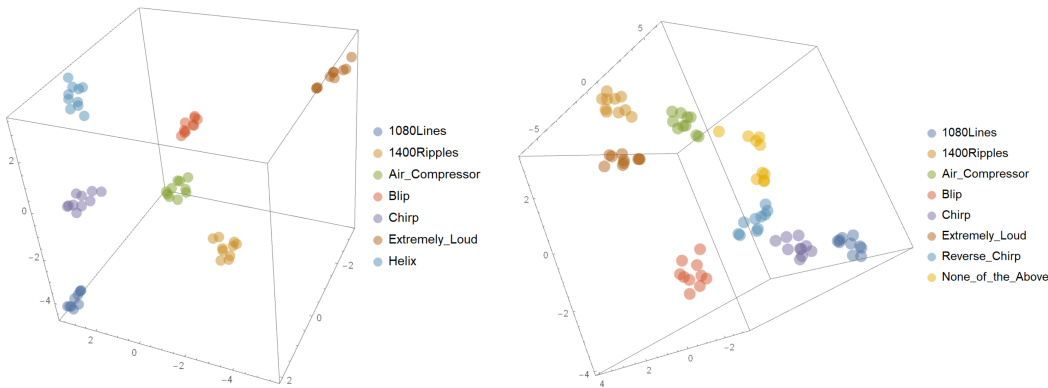


Figure 2: Location of different classes of glitches from the test set after applying the CNN feature-extractor from our fine-tuned InceptionV3 model. The t-SNE [32] algorithm was used to reduce the dimension to a 3D vector. Note that each type of glitch forms a cluster, and their relative positions depend on their morphology. Outliers may be inspected closely to verify their labels and decide whether they should belong to a new class. A new class called Reverse\_Chirp was added. It can be seen that the CNN feature-extractor maps this class (which was not shown during training) to a unique cluster. Furthermore, this cluster is located near the Chirp class and the None\_of\_the\_Above class, which means that the relative positions of the glitches in this feature-space is meaningful. Note that glitches in the None\_of\_the\_Above are also grouped into smaller clusters.

we achieved perfect precision and recall on 7 classes: 1080Lines, 1400Ripples, Extremely\_Loud, Helix, Paired\_Doves, Scratchy, and Violin\_Mode. Both ResNet50 and InceptionV3 achieved the highest accuracy of 98.84% on the test set despite being trained independently via different methods on different splits of the data. Both models obtained 100.00% accuracy when considering the top-5 predictions, which implies that given any input, the true class can be narrowed down to within 5 classes with 100.00% confidence. This is particularly useful, since the true class of a glitch is often ambiguous to even human experts.

We found that the trained CNNs may also be used as good feature extractors for finding new categories of glitches from unlabeled data in an unsupervised or semi-supervised manner. This method can be used to identify many more categories of noise transients and estimate at what times new types of glitches with similar morphologies start occurring. This may also be used to correct mislabeled glitches in the original dataset used for training/testing by searching for anomalies in the feature-space. For any of our models, removing the final softmax and fully-connected layer near the output produces a CNN that maps any input image to a vector of real numbers which encode useful information distinguishing different classes of glitches. In this high-dimensional space, glitches having similar morphology will be clustered together. Therefore, when new types of glitches appear, which are classified as None\_of\_the\_Above by our CNN model, they may be mapped to vectors using these truncated CNN feature extractors (see Figure 2) and new clusters (classes) can be found.




Figure 3: Confusion matrix for InceptionV3. The accuracy is 98.84% for classification on the test set. The *Chirp* glitch class (which is also the class that would contain true GW signals from mergers of black hole or neutron star binaries) and the *Paired\_Doves* class (the smallest class, which only had 24 training examples) was identified with perfect precision and recall, i.e., perfect accuracy.

## 4 Conclusion

In this article, we have developed state-of-the-art CNNs for LIGO glitch classification using the Gravity Spy dataset. We have shown that by applying transfer learning from ImageNet, we can obtain excellent results with small training datasets and achieve significantly better the accuracy compared to CNNs trained from scratch only on glitch spectrograms. Furthermore, the training time is significantly reduced with our approach by several orders of magnitude, and the effort required to design CNN models and optimize their hyper-parameters can be eliminated. The algorithms we have introduced in this paper may be used to classify new time-series data in the Gravity Spy project, as well as data streams in real-time from future LIGO and Virgo observing runs as well as KAGRA [33] and LIGO-India [34], as they come online in the next few years. The transfer learning method also allows us to use the fine-tuned CNNs as feature extractors for clustering algorithms to find new classes of glitches and signals in an unsupervised manner or to label them rapidly in a semi-supervised manner. We expect our methods for glitch classification and clustering may help in finding the instrumental or environmental sources of many classes of glitches whose origins remain unknown, prevent false detection, and enhance the quality of data from gravitational wave detectors thus enabling new scientific discoveries. Furthermore, we anticipate that these techniques may also be useful in general for detecting, classifying, and clustering anomalies in other disciplines.



## Acknowledgements

We thank Gabrielle Allen, Ed Seidel, Scott Coughlin, Vicky Kalogera, Aggelos Katsaggelos, Joshua Smith, Kai Staats, Sara Bahaadini, and Michael Zevin for productive interactions. We thank Kai Staats, Laura Nuttall, the [Gravity Spy](#) team, [NCSA Gravity Group](#), and many others for reviewing this article and providing feedback. We are grateful to NVIDIA for supporting this research by donating four P100 GPUs, to Wolfram Research for offering several [Wolfram Language](#) (*Mathematica*) licenses used for this work, and to Vlad Kindratenko for providing dedicated access to a high-performance machine at the Innovative Systems Lab at NCSA. We acknowledge the Gravity Spy project and the citizen scientists who participated in it for processing and labeling the raw data from LIGO. We also acknowledge the LIGO collaboration for the use of computational resources and for the feedback from the CBC, DetChar, and MLA working groups.

## References

- [1] The LIGO Scientific Collaboration, J. Aasi, et al. Advanced LIGO. *Classical and Quantum Gravity*, 32(7):074001, April 2015.
- [2] F. Acernese et al. for the Virgo Collaboration. Advanced Virgo: a second-generation interferometric gravitational wave detector. *Classical and Quantum Gravity*, 32(2):024001, January 2015.
- [3] The LIGO Scientific Collaboration and The Virgo Collaboration. GW150914: The Advanced LIGO Detectors in the Era of First Discoveries. *ArXiv e-prints*, February 2016.
- [4] B. P. et al Abbott. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, Feb 2016.
- [5] B. P. et al Abbott. Gw151226: Observation of gravitational waves from a 22-solar-mass binary black hole coalescence. *Phys. Rev. Lett.*, 116:241103, Jun 2016.
- [6] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, and et al. GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2. *Physical Review Letters*, 118:221101, Jun 2017.
- [7] B. P. et al. Abbott. Gw170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Phys. Rev. Lett.*, 119:141101, Oct 2017.
- [8] B. P. et al. Abbott. Gw170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys. Rev. Lett.*, 119:161101, Oct 2017.
- [9] LIGO Scientific Collaboration, Virgo Collaboration, F. GBM, INTEGRAL, IceCube Collaboration, AstroSat Cadmium Zinc Telluride Imager Team, IPN Collaboration, The Insight-Hxmt Collaboration, ANTARES Collaboration, The Swift Collaboration, AGILE Team, The 1M2H Team, The Dark Energy Camera GW-EM Collaboration, the DES Collaboration, The DLT40 Collaboration, GRAWITA, :, GRAvitational Wave Inaf TeAm, The Fermi Large Area Telescope Collaboration, ATCA, :, A. Telescope Compact Array, ASKAP, :, A. SKA Pathfinder, Las Cumbres Observatory Group, OzGrav, DWF, AST3, CAASTRO Collaborations, The VIN-ROUGE Collaboration, MASTER Collaboration, J-GEM, GROWTH, JAGWAR, C. NRAO, TTU-NRAO, et al. Multi-messenger Observations of a Binary Neutron Star Merger. *ArXiv e-prints*, October 2017.
- [10] LIGO Scientific Collaboration, Virgo Collaboration, J. Aasi, J. Abadie, B. P. Abbott, R. Abbott, T. D. Abbott, M. Abernathy, T. Accadia, F. Acernese, et al. Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *ArXiv e-prints*, April 2013.
- [11] B. S. Sathyaprakash and B. F. Schutz. Physics, Astrophysics and Cosmology with Gravitational Waves. *Living Reviews in Relativity*, 12:2, March 2009.
- [12] N. J. Cornish and T. B. Littenberg. Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches. *Classical and Quantum Gravity*, 32(13):135012, July 2015.
- [13] J. Powell, D. Trifirò, E. Cuoco, I. S. Heng, and M. Cavaglià. Classification methods for noise transients in advanced gravitational-wave detectors. *Classical and Quantum Gravity*, 32(21):215012, November 2015.

- [14] J. Powell, A. Torres-Forné, R. Lynch, D. Trifirò, E. Cuoco, M. Cavaglià, I. S. Heng, and J. A. Font. Classification methods for noise transients in advanced gravitational-wave detectors II: performance tests on Advanced LIGO data. *ArXiv e-prints*, September 2016.
- [15] N. Mukund, S. Abraham, S. Kandhasamy, S. Mitra, and N. S. Philip. Transient classification in ligo data using difference boosting neural network. *Phys. Rev. D*, 95:104059, May 2017.
- [16] D. George and E. A. Huerta. Deep Neural Networks to Enable Real-time Multimessenger Astrophysics. *ArXiv e-prints*, December 2017.
- [17] D. George and E. A. Huerta. Deep Learning for Real-time Gravitational Wave Detection and Parameter Estimation: Results with Advanced LIGO Data. *ArXiv e-prints*, November 2017.
- [18] D. George, H. Shen, and E. A. Huerta. Deep Transfer Learning: A new deep learning glitch classification method for advanced LIGO. *ArXiv e-prints*, June 2017.
- [19] The LIGO Scientific Collaboration and the Virgo Collaboration. Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914. *ArXiv e-prints*, February 2016.
- [20] The LIGO Scientific Collaboration and B. P. Abbott. Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914. *ArXiv e-prints*, February 2016.
- [21] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. Katsaggelos, S. Larson, T. K. Lee, C. Lintott, T. Littenberg, A. Lundgren, C. Oesterlund, J. Smith, L. Trouille, and V. Kalogera. Gravity Spy: Integrating Advanced LIGO Detector Characterization, Machine Learning, and Citizen Science. *Classical and Quantum Gravity*, 34(11), February 2017.
- [22] D. George and E. A. Huerta. Deep Learning for Real-time Gravitational Wave Detection and Parameter Estimation with Advanced LIGO Data. *ArXiv e-prints*, November 2017.
- [23] H. Shen, D. George, E. A. Huerta, and Z. Zhao. Denoising Gravitational Waves using Deep Learning with Recurrent Denoising Autoencoders. *ArXiv e-prints*, November 2017.
- [24] Yann Lecun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 5 2015.
- [25] S. Bahaadini, N. Rohani, S. Coughlin, M. Zevin, V. Kalogera, and A. K. Katsaggelos. Deep Multi-view Models for Glitch Classification. *International Conference on Acoustics, Speech and Signal Processing*, April 2017.
- [26] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS'14*, pages 3320–3328, Cambridge, MA, USA, 2014. MIT Press.
- [27] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. CNN Features off-the-shelf: an Astounding Baseline for Recognition. *ArXiv e-prints*, March 2014.
- [28] S. Christian, L. Wei, Jia. Yangqing, S. Pierre, R. Scott, A. Dragomir, E. Dumitru, V. Vincent, and R. Andrew. Going Deeper with Convolutions. *ArXiv e-prints*, September 2014.
- [29] H. Kaiming, Z. Xiangyu, R. Shaoqing, and S. Jian. Deep Residual Learning for Image Recognition. *ArXiv e-prints*, December 2015.
- [30] S. Karen and Z. Andrew. Very Deep Convolutional Networks for Large-Scale Image Recognition. *ArXiv e-prints*, September 2014.
- [31] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. *CVRP*, 2009.
- [32] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.
- [33] E. Hirose, T. Sekiguchi, R. Kumar, R. Takahashi, and the KAGRA Collaboration. Update on the development of cryogenic sapphire mirrors and their seismic attenuation system for KAGRA. *Classical and Quantum Gravity*, 31(22):224004, November 2014.
- [34] C. S. Unnikrishnan. IndIGO and Ligo-India Scope and Plans for Gravitational Wave Research and Precision Metrology in India. *International Journal of Modern Physics D*, 22:41010, January 2013.