# Advancing Glitch Classification in Gravity Spy: Multi-view Fusion with Attention-based Machine Learning for Advanced LIGO's Fourth Observing Run

Yunan Wu<sup>a,\*</sup>, Michael Zevin<sup>b,e</sup>, Christopher P. L. Berry<sup>c</sup>, Kevin Crowston<sup>d</sup>, Carsten Østerlund<sup>d</sup>, Zoheyr Doctor<sup>e</sup>, Sharan Banagiri<sup>e</sup>, Corey B. Jackson<sup>f</sup>, Vicky Kalogera<sup>e,g</sup>, Aggelos K. Katsaggelos<sup>a,e</sup>

#### Abstract

The first successful detection of gravitational waves by ground-based observatories, such as the Laser Interferometer Gravitational-Wave Observatory (LIGO), marked a revolutionary breakthrough in our comprehension of the Universe. However, due to the unprecedented sensitivity required to make such observations, gravitational-wave detectors also capture disruptive noise sources called glitches, potentially masking or appearing as gravitational-wave signals themselves. To address this problem, a community-science project, Gravity Spy, incorporates human insight and machine learning to classify glitches in LIGO data. The machine learning classifier, integrated

Email address: yunanwu2020@u.northwestern.edu (Yunan Wu)

<sup>&</sup>lt;sup>a</sup>The Department of Electrical Computer Engineering, Northwestern University, 2145 Sheridan Road, Evanston, 60208, IL, USA

<sup>&</sup>lt;sup>b</sup> The Adler Planetarium, 1300 South DuSable Lake Shore Drive, Chicago, 60605, IL, USA

<sup>&</sup>lt;sup>c</sup>SUPA, School of Physics and Astronomy, University of Glasgow, Kelvin Building, University Ave, Glasgow, 8QQ, G12, UK

<sup>&</sup>lt;sup>d</sup>School of Information Studies, Syracuse University, 343 Hinds Hall, Syracuse, 13210, NY, USA

 <sup>&</sup>lt;sup>e</sup>Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, 1800 Sherman Ave, Evanston, 60201, IL, USA
<sup>f</sup>Information School, University of Wisconsin-Madison, 600 N Park Street, Madison, 53706. WI. USA

<sup>&</sup>lt;sup>g</sup>Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, 60208, IL, USA

<sup>\*</sup>Corresponding author.

into the project since 2017, has evolved over time to accommodate increasing numbers of glitch classes. Despite its success, limitations have arisen in the ongoing LIGO fourth observing run (O4) due to its architecture's simplicity, which led to poor generalization and inability to handle multi-time window inputs effectively. We propose an advanced classifier for O4 glitches. Our contributions include evaluating fusion strategies for multi-time window inputs, using label smoothing to counter noisy labels, and enhancing interpretability through attention module-generated weights. This development seeks to enhance glitch classification, aiding in the ongoing exploration of gravitational-wave phenomena.

Keywords: LIGO, Glitch Classification, Gravity Spy, O4, Machine Learning, Attention

#### 1. Introduction

The first discovery of gravitational waves, a pivotal element in Einstein's theory of general relativity [1], has opened up an entirely new window in the cosmos. The Laser Interferometer Gravitational-Wave Observatory (LIGO) [2] achieved groundbreaking success in detecting these ripples in spacetime for the first time in September 2015 [3], fundamentally transforming our understanding of the Universe. Later on, Virgo [4] made its first detection of gravitational waves in August 2017 [5, 6] and KAGRA initiated its gravitational waves observations in April 2020 [7, 8]. However, achieving these observations demands exceptionally sensitive and intricate detectors to be able to measure minuscule fluctuations in spacetime [9]. This heightened sensitivity, in turn, leads to the detection of diverse sources of noise, with the potential to obscure or mimic authentic gravitational-wave signals [10, 11]. Among these noises, bursts of non-Gaussian noise caused by environmental or instrumental factors, known as glitches, are particularly disruptive to measuring gravitational waves. The origins of many glitches remain largely unknown [12, 13]. The process of characterizing and eliminating these glitches is essential to enhancing the quality of the detection system, thereby increasing the likelihood of analyzing gravitational waves.

In pursuit of this objective, *Gravity Spy* [14, 15], a community-science project, has been launched to leverage the power of both volunteers and machine learning to classify glitches in the LIGO time-series data. On one front, the project engages human expertise in identifying and categorizing various

glitches collected by LIGO. These labeled glitches then serve as training data to refine machine learning classifiers, enhancing the accuracy and efficiency of glitch identification. Conversely, the classifiers provide guidance to volunteers, assisting them in identifying existing or new potential glitch classes [16]. Gravity Spy is implemented on the Zooniverse platform [17] on the community-science side, and it has demonstrated significant success by attracting over 30 thousand volunteers making more than 7 million glitch classifications. The outputs are actively used by LIGO. One critical component contributing to this significant achievement is the integration of the machine learning classifier.

The machine learning classifier has been seamlessly integrated into the Gravity Spy pipeline for over five years, largely maintaining its original architecture. Periodically, fine-tuning of its final layer has been performed to enhance its adaptability to diversely predicted glitch classes [18]. The input for the classifier remains unchanged, featuring a structure that juxtaposes four time window spectrograms (0.5 s, 1.0 s, 2.0 s and 4.0 s) of a single glitch into one merged image. The first version of the machine learning classifier was introduced in the initial Gravity Spy paper [14] and integrated into the Gravity Spy pipeline after the completion of the first observing run (O1). Designed for a 20-class classification setup, this classifier consisted of two convolutional layers and two fully-connected layers to extract image features. Each convolutional layer was followed by a max pooling layer to reduce the dimensionality of the features.

As the project progressed, during the second observing run (O2), Gravity Spy volunteers identified new glitch classes and two of them, i.e., 1080 Lines and 1400 Ripples, were added into the classifier. Therefore, the classifier was optimized by incorporating two additional convolutional layers and max pooling layers before the existing fully-connected layers, which facilitated the extraction of more discriminative features and adapted the classifier to a 22-class classification task [19].

During the third observing run (O3), two new glitch classes, known as Low Frequency Blip and Fast Scattering, were identified through the collaborative efforts of volunteers and LIGO experts [20]. Alongside this, the decision was made to remove the None of the Above class from the training dataset. The None of the Above class was initially aimed at empowering our volunteers

<sup>&</sup>lt;sup>1</sup>gravityspy.org

to identify new glitch classes. However, its inclusion confuses the model, especially given morphological similarities with other glitches. As a result, removing this class allows the identification of new glitches more effectively with low predicted confidence [20, 21]. This led to further retraining of the classifier, enabling it to effectively accommodate a 23-class classification task [20].

Currently, the fourth observing run (O4) of the LIGO detector network is in process [22]. With the anticipation of discovering new glitch classes during O4 and subsequent runs, the current Gravity Spy classifier faces major limitations [23, 24]. First, in the glitch classification task, unlike other image processing tasks, location plays a crucial role. For example, distinguishing between a line at the top and a line at the bottom of the image is essential as they represent different glitches at different frequencies. However, the classifier's relatively straightforward architecture, comprised mainly of a combination of convolutional and fully-connected layers, hinders its capacity to capture unique features in individual glitch classes [25]. Second, due to its shallow layer structure, the classifier requires significant resizing of input image dimensions, resulting in the loss of valuable information. Moreover, the classifier uses a single model to extract features from merged image inputs across different time windows, which further constrains the ability to uncover cross-time window correlations [26]. Recognizing these challenges, it becomes apparent that the demand to analyze more data and diverse glitch types exceeds the current model's capabilities.

Additionally, the current classifier grapples with a confidence overfitting issue, displaying excessive certainty in glitch predictions even when errors are present. This tendency can potentially mislead both volunteers and LIGO experts, posing a risk to the efficiency and reliability of glitch characterization. Therefore, there is a compelling need for an advanced classifier architecture with robust generalization capabilities in the context of O4.

In this study, we develop a novel machine learning classifier for glitch classification in the ongoing O4 run to address these challenges. The proposed classifier employs an attention-based multi-view fusion strategy to capture cross-time window correlations across the four time windows. The incorporation of a regularizer and label smoothing techniques into the loss function improves generalization and mitigates the confidence overfitting issue. The contributions of this paper can be summarized as follows: 1) We comprehensively compare three fusion strategies (early fusion, late fusion, and intermediate fusion) in scenarios with multiple time windows as inputs. 2) We

apply label smoothing to mitigate the impact of noisy training labels. 3) We introduce attention modules in the classifier, enhancing transparency in the glitch decision-making process by identifying specific window time that predominantly influences the glitch classification decision — an exploration undertaken for the first time.

The paper is organized as follows: in Section 2, we provide an overview of the Gravity Spy dataset, introduce three fusion strategies, outline the architecture of the proposed classifier, and elaborate on the conducted statistical analysis. Moving on to Section 3, we present the outcomes obtained from various ablation studies. In Section 4, we delve into a comprehensive discussion, addressing both the strengths and limitations of this study. We conclude in Section 5.

## 2. Materials and Methods

In this section, we provide an overview of the Gravity Spy dataset, covering its data generation process, image preprocessing, and the methodology for data splitting during both training and testing phases. Following this, we delve into the multi-view fusion strategy, discussing its early, intermediate, and late fusion components. Last, we propose a novel classifier, outlining the model architecture, attention modules, and the application of a label smoothing technique.

## 2.1. Gravity Spy Dataset

The Gravity Spy dataset comprises time-frequency spectrograms of glitches. The glitches are initially identified by the Omicron trigger pipeline [27]. To effectively capture glitches that pose significant challenges to gravitational-wave searches, this pipeline employs strict criteria, identifying only noise events with a signal-to-noise ratio exceeding 7.5 and a peak frequency falling within the range of 10 Hz to 2048 Hz. Subsequently, the Q-transform process [28, 29] is employed to convert the LIGO strain data into time-frequency spectrograms, known as Omega scans. This involves segmenting the strain time series into shorter intervals and applying the discrete Fourier transform to extract the frequency content. Gravity Spy generates four spectrograms for each glitch triggered, using distinct time windows: 0.5 s, 1.0 s, 2.0 s, and 4.0 s. The incorporation of multiple time windows serves a crucial purpose, allowing both volunteers and machine learning models to analyze the morphologies of glitches occurring at different characteristic timescales. For

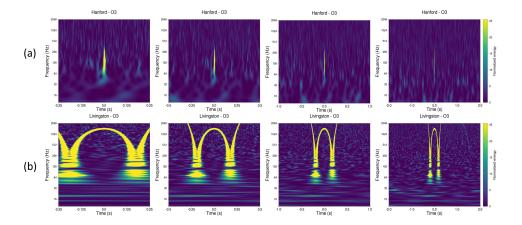


Figure 1: Spectrogram examples of two glitches with four time windows (0.5 s, 1.0 s, 2.0 s and 4.0 s). (a) Blip, characterized by its short duration. (b) Whistle, characterized by its long duration. These examples illustrate how different time windows enable the analysis of glitches occurring at various temporal scales, providing valuable insights into the morphological characteristics of each glitch type.

instance, this enables the examination of short (i.e., Blip) and long duration (i.e., Whistle) noise events, as depicted in Fig. 1.

The training and testing datasets employed in this study are sourced from Zevin et al. [30], focusing exclusively on data from the Hanford and Livingston detectors during the third observing run (O3) of Advanced LIGO. This classification task has in total 23 glitch classes. To ensure consistency, the names associated with each class and the typical morphology of glitches belonging to each class were determined through a collaborative effort between LIGO scientists and citizen scientists. Examples representing each of the 23 glitch classes are illustrated in Fig. 2 and more details are provided in the O1-O3 data release paper [21].

For the purposes of this study, we consider the ground truth for glitches in our dataset to be established through a collaborative effort involving volunteers and the previously implemented machine learning classifier [14]. However, this ground truth may exhibit a potential noisy-label issue, that is a few glitches were misclassified by the classifier and some volunteers. To strengthen the confidence in the ground truth used for our new classification model, we adopt a filtering criterion to create the training and testing dataset. Specifically, we only use glitches with consensus confidence scores

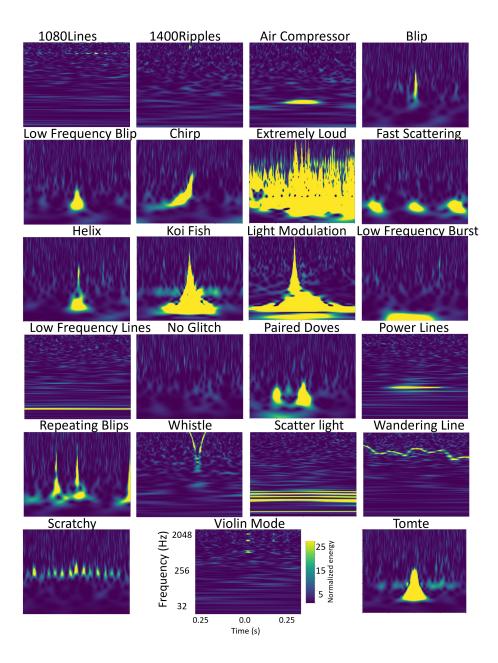


Figure 2: Spectrogram examples of each of the 23 glitch classes within the Gravity Spy dataset. All glitches share the same axes and scale bar with Violin Mode. These examples provide one possible morphological representation for each class, while each class can exhibit various shapes and patterns [21].

from both machine learning and volunteers greater than 0.9. As a consequence, our new model is trained on glitches with relatively reliable labels, effectively reducing confusion during the training process and shielding the model from potentially mislabeled ground truth. Improved training sets for O4 data will be the focus of future studies.

Following the filtering criteria, the Gravity Spy dataset is divided into three distinct sets. The training set consists of 8439 glitches and the validation set consists of 1687 glitches. Additionally, a separate set of 3538 glitches is reserved for testing purposes. Each set is carefully curated to ensure no overlapping glitch subjects, thereby preventing data leakage between those three phases. A detailed breakdown of the data splitting is presented in Table 1.

Table 1: The distribution of each class in the Gravity Spy dataset across the training, validation, and testing sets.

Class	Train	Validation	Testing	Class	Train	Validation	Testing
1080Lines	184	49	225	Low Frequency Lines	381	91	225
1400Ripples	257	64	224	No Glitch	311	78	225
Air Compressor	293	57	225	Paired Doves	124	31	146
Blip	312	84	218	Power Line	354	83	224
Low Frequency Blip	337	103	222	Repeating Blips	411	90	224
Chirp	38	10	17	Scattered Light	521	128	224
Extremely Loud	327	85	224	Scratchy	178	36	225
Fast Scattering	515	128	224	Tomte	408	90	225
Helix	46	10	71	Violin Mode	210	56	225
Koi Fish	338	93	224	Wandering Line	25	8	25
Light Modulation	127	28	159	Whistle	462	129	224
Low Frequency Burst	378	106	223				

The glitches from all three cohorts undergo the same pre-processing steps. Initially, each spectrogram's original dimension size is  $600 \times 800 \times 3$ . We first crop the spectrogram with a bounding box defined as [top-left vertical:60, top-left horizontal:100, height:580, width:675] to retain only the intensity region of interest and eliminate irrelevant information, such as axes and scaling maps. Next, all cropped glitches are resized to a standardized size of  $448 \times 448 \times 3$  to facilitate uniformity during model training. Finally, minmax normalization is applied, transforming the pixel intensities into a range of 0 to 1, ensuring that all glitches are equally scaled and enhancing the model's ability to effectively learn from the data.

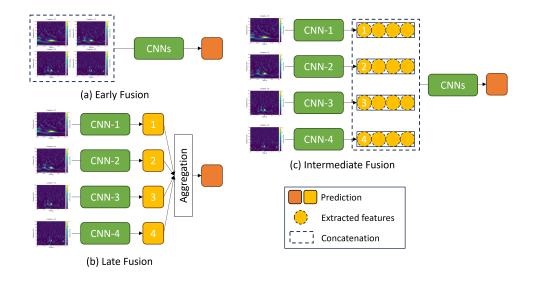


Figure 3: Illustration of three fusion strategies. (a) Early fusion; (b) Late fusion, and (c) Intermediate fusion

## 2.2. Multi-view Fusion Strategy

As Gravity Spy generates a dataset of glitches with four different durations, it is crucial to have a method for incorporating glitch information of varying characteristic timescales. Multi-view fusion, a powerful strategy in machine learning, combines information from multiple sources to enhance the representation and performance of a model. Each view contributes unique and complementary insights, leading to improved accuracy, robustness, and generalization of the model [26]. As mentioned above, our Gravity Spy project serves as an example where a single image with four time windows enables models to extract features across diverse timescales. Typically, multi-view fusion can employ one of three fusion architectures with different strategies to combine data from various view images. We review those below: the performance of these three fusion strategies will be compared in the result section.

#### 2.2.1. Early Fusion

Early fusion takes place at the input level, where features from each view are combined in the input layer of the network before being fed into one single model. By merging data from multiple views at this stage, the

model has access to a more comprehensive representation of the data. This approach was employed in previous Gravity Spy classifiers [21, 20, 14, 31], which constructed a  $2h \times 2w$  matrix by juxtaposing four  $h \times w$  view images, as illustrated in Fig. 3a.

Although it offers several advantages, it has some limitations that may impact its practicality, particularly for large-sized images. In previous classifiers [21, 20, 14, 31], the resizing of each view to  $140 \times 170$  was required to fit within the model's limitations, leading to a loss of valuable information. Additionally, early fusion assumes equal contributions from all modalities to the task, which may not hold true in this specific project. For instance, Repeating Blips might only appear in the 4 s time window view and not in others, necessitating the model to assign greater weight to the 4 s window for accurate classification.

#### 2.2.2. Late Fusion

Late fusion is a decision-level data fusion technique, where predictions or outputs from individual models trained on different views are combined at the final stage, as shown in Fig. 3b. This fusion process often involves aggregation methods, such as majority voting or averaging, to make the final decision. The advantage of late fusion is its flexibility in utilizing different models for each view, enabling the selection of the most suitable model for each input. Moreover, it handles missing views by excluding the corresponding model during the fusion process, ensuring model evaluation even when certain data sources are unavailable (though that was not in the case in this application). However, one drawback of late fusion is its potential to overlook cross-modal relationships, as fusion occurs at a higher decision level rather than at the earlier processing stages. Consequently, in our case, this might lead to suboptimal exploitation of valuable complementarity between different time-window views.

# 2.2.3. Intermediate Fusion

Intermediate fusion, also known as feature-level fusion, allows features from each view to be processed separately through different branches of the network, and then combines the resulting feature maps at an intermediate layer, as shown in Fig. 3c. This technique effectively addresses the limitations of both early fusion and late fusion. Specifically, by leveraging the complementarity of different views and capturing cross-modal relationships at a stage where features are still semantically meaningful, intermediate fusion

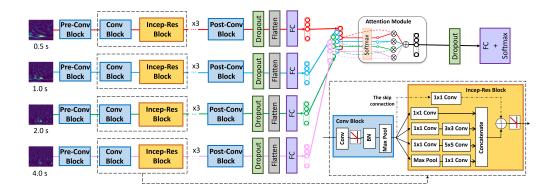


Figure 4: Architecture of the proposed machine learning classifier used in O4. Further elaboration on the hyperparameters can be found in Section 2.3.

enables the model to learn complex representations and potentially achieve better performance. One general issue with intermediate fusion is the handling of missing views. As the model relies on inputs from all views, the absence of one or more views can lead to the loss of crucial information, potentially affecting the model's decision-making process. This problem is not a concern for this project, as each glitch always contains complete data from all four time windows during pre-processing.

## 2.3. Methods

The proposed classifier distinguishes itself from glitch classification through several key optimizations, addressing the limitations of previous models [21, 20, 14, 31]. First, it adopts intermediate fusion, effectively combining four views to enhance its representation power by leveraging their complementary information. Second, to handle a larger input dimension, a deeper network is adopted, featuring the inception residual block [32] to mitigate the vanishing gradient problem and to extract more discriminative features. Additionally, an attention module is introduced at a later stage, providing valuable insights into the model's focus on specific views during decision-making. Last, the incorporation of smooth labeling during training effectively addresses the noisy label issue. The architecture of the proposed classifier is shown in Fig. 4, and detailed explanations of these optimizations are presented next.

## 2.3.1. Inception Residual Block

The inception residual block combines the ideas of the inception module and the residual block proposed in Szegedy et al. [32]. As shown in Fig. 4, it combines these concepts by applying the inception module within a residual block structure. Specifically, the block consists of four parallel convolutional branches with filters of different sizes  $(1 \times 1, 1 \times 1, 3 \times 3, 5 \times 5)$ . The outputs of these branches are then concatenated, and the resulting feature map is combined with the original input after a  $1 \times 1$  convolutional layer using a skip connection, followed by a rectified linear unit (ReLU) activation function. Particularly, the use of parallel convolutional filters of different sizes allows the block to capture features at multiple scales, improving the model's ability to recognize patterns of varying complexities [33]. Moreover, the incorporation of skip connections facilitates the training of deep networks, effectively mitigating the vanishing gradient problem [34].

#### 2.3.2. Attention Module

To assess the varying importance of each view in the final glitch classification, we introduce an attention module [35]  $\mathcal{M}_{\mathrm{Att}}$  to the latent vectors of four inputs, denoted as  $z_i \in \mathbf{R}^{1 \times M}, i = \{1, 2, 3, 4\}$ , extracted after the fully-connected layer. The attention weight  $\alpha_i \in \mathbf{R}, i = \{1, 2, 3, 4\}$  calculated for each latent vector  $z_i$  quantifies the significance of each view in making the glitch classification, with all  $\alpha_i$ 's summing up to 1. The attention module is defined as:

$$\mathcal{M}_{\text{Att}} = \sum_{i=1}^{4} \alpha_i z_i,\tag{1}$$

where the attention weight  $\alpha_i$  is computed as

$$\alpha_i = \frac{\exp\{w^T[\tanh(Vz_i^T) \odot \sigma(Uz_i^T)]\}}{\sum_{i=1}^4 \exp\{w^T[\tanh(Vz_i^T) \odot \sigma(Uz_i^T)]\}}.$$
 (2)

The trainable parameters  $V \in \mathbf{R}^{L \times M}$  and  $w \in \mathbf{R}^{L \times 1}$  are associated with the attention mechanism. Here, M represents the dimension of latent vectors (set to M=128 in this paper) while L signifies the one dimension of the matrices for transforming the latent vectors (set to L=64). The V matrix is instrumental in identifying similarities among different views. The nonlinear hyperbolic tangent function  $\tanh(\cdot)$  includes both positive and negative values in the gradient flow, but the concern is that it is almost linear in the

[-1,1] range. To handle complex data patterns effectively, we leverage the gated attention mechanism [35], which combines  $\tanh(\cdot)$  with the elementwise sigmoid function  $\sigma(\cdot)$ . The sigmoid function helps adjust the relevance of information, while the element-wise multiplication  $\odot$  facilitates the attention process. The gated attention mechanism introduces additional parameters  $U \in \mathbf{R}^{L \times M}$  to serve as control gates, mitigating the linearity issue present in  $\tanh(\cdot)$ . The attention module, widely used in multiple-instance learning tasks [35, 36, 37, 38], excels in identifying key instances within a bag, i.e., a collection of instances.

In this paper, we assess the impact of the attention module on the model's performance through comparative experiments conducted with and without the incorporation of attention modules.

## 2.3.3. Label Smoothing

As mentioned in Sec. 2.1, the ground truth labels for the Gravity Spy dataset are generated by the previous classifier and volunteers. Despite filtering the dataset to only use high-confidence predictions, mislabeled glitches can still be present, potentially confusing the model. To address this problem, label smoothing [39], a valuable regularization technique in machine learning, is applied during model training to enhance generalization and prevent overconfidence in predictions, especially when dealing with noisy or uncertain data. This technique introduces a small amount of uncertainty into the training labels. Instead of assigning a hard 1 to the correct class and 0 to other classes, label smoothing assigns a smoothed value, denoted by q, slightly less than 1 to the correct class while distributing the remaining confidence uniformly across the other classes. Mathematically, it can be defined as:

$$q = (1 - \beta)y + \beta u, (3)$$

where y represents the true label (encoded as a one-hot vector), u is a uniform distribution over the classes, and  $\beta$  is a smoothing parameter between 0 and 1.

In addition, to address the issue of class imbalance (as shown in Table 1), we apply different weights  $\gamma_i$  to each class  $i = \{1, 2, ...C\}$  in the loss function during training. These weights are inversely proportional to the number of instances in each glitch class and sum up to 1, ensuring a balanced learning

process. Therefore, the loss function of the proposed model is defined as

$$\mathcal{L}(x,y) = -\frac{1}{C} \sum_{i=1}^{C} \gamma_i q_i \log(\hat{y}_i), \tag{4}$$

where  $\hat{y}$  represents the predicted probability distribution over classes and C is the number of glitch classes. By adjusting the smoothing parameter, the level of smoothing applied to the labels can be controlled. For example, setting  $\beta=0$  (i.e., no smoothing) leads to one-hot encodings and restores the standard categorical cross-entropy loss, while  $\beta=1$  (i.e., completely smooth) results in a uniform distribution over all classes.

In this paper, the impact of different smoothing levels on the model's performance is also evaluated through different experiments on the value of  $\beta$ .

# 2.3.4. Architecture of the Classifier

The overall architecture of the proposed classifier is shown in Fig. 4. Each of the four branches shares a common architecture. Within each branch, a pre-convolutional block is followed by three instances of a combination of a convolutional block and an inception residual block (Incep-Res Block) [32], along with a post-convolutional block. These stages effectively extract features from each glitch view. Specifically, all convolutional blocks (i.e., Pre-Conv., Conv and post-Conv blocks) consist of a convolutional layer, a ReLu activation function, and batch normalization (BN) layer and a max pooling (Max Pool) layer. Subsequently, a flatten layer and a fully-connected (FC) layer process and reduce the dimensionality of the feature vectors. The attention module [35] then takes these vectors as input, generating a single weighted averaged feature vector, which is then fed into the final FC layer. The softmax activation function yields the probabilities of the 23 glitch classes. To enhance the model's generalization and prevent overfitting, we incorporate an L2 regularizer [40] for each convolutional layer and introduce dropout [41] after each FC layer. For optimization, we use the Adam optimizer [42], with the learning rate starting at  $2 \times 10^{-5}$ . This choice of optimizer helps to efficiently update the model's parameters during training. To address the potential issue of confidence overfitting, where the model may become overconfident in its predictions on the training data, we employ early stopping with a patience of 10, i.e., if the model doesn't improve its accuracy on the validation set within 10 rounds of training, we stop training to prevent potential overconfidence.

## 2.4. Statistical Analysis

To ensure the robustness of the model, five independent runs for all experiments are conducted. For each experiment, we calculate the overall accuracy, precision, F1 score, and the area under the receiver operating characteristic curve (AUC) on the held-out test set, where the mean and standard deviation values for each metric are reported. To assess the significance of the differences between different model results, we perform a paired t-test [43], considering the comparison statistically significant if the p-value is less than 0.05 (p < 0.05). Furthermore, to calibrate the performance of the model, a reliability plot [44] is employed, offering a visual representation of the model's calibration across varying confidence levels. Specifically, this plot juxtaposes predicted probabilities against empirical frequencies of positive outcomes, providing insights into the model's accuracy in conveying uncertainty. Additionally, calibration accuracy is quantified through the expected calibration error (ECE) [45], a statistical metric that measures the average deviation between predicted probabilities and the true likelihood of event occurrence.

## 3. Results

In this section, we present an overview of our results. We begin by outlining the design of all experiments conducted, providing a comprehensive understanding of our methods. Next, we compare the classifier's results under various fusion strategies, evaluate the impact of different smoothing levels, analyze the effects of the gated attention mechanism, and, last, perform a comparative assessment between the proposed model and the previous model implemented in the Gravity Spy pipeline.

# 3.1. Experimental Design

In this study, we conduct various ablation studies to assess the model's performance. First, the effectiveness of different fusion strategies is compared, as shown in Fig. 3. Specifically, the early fusion involves merging four view images and resizing the merged input to a size of  $448 \times 448 \times 3$  before feeding it into the model. The late fusion uses each branch to individually produce probabilities for each class, and the final classification is determined by averaging the probabilities obtained from the four view images. The proposed classifier employs intermediate fusion. Second, we evaluate different values of the smoothing parameter  $\beta$  between 0 and 1 with an increment of 0.1, allowing us to identify the optimal value for  $\beta$ . Third, we compare

the proposed classifier with the attention module and a variant without the attention module (i.e., max pooling and mean pooling) to demonstrate the effectiveness of the attention mechanism. Fourth, we evaluate the performance of our proposed classifier by comparing it with the previous Gravity Spy model architecture [20, 21]. Both models are trained from scratch using the same training dataset, and their performance is assessed on the same test set. Finally, to better understand the learned feature representations, we employ t-distributed stochastic neighbor embedding (t-SNE) plots to reduce the dimensionality of feature vectors obtained after the attention module from 128 to 2, offering insightful visualizations for model comparisons in a two-dimensional space. The experiments and analyses are conducted using TensorFlow 2.11 on two GPUs (NVIDIA Quadro RTX 8000) and scikit-learn, respectively, in Python 3.8. The codes are available via the GitHub of Gravity Spy.

# 3.2. Comparison of Three Fusion Strategies

Table 2 demonstrates the model performance of three fusion strategies on the same test set. The proposed intermediate fusion model achieves the best result with an overall accuracy of 0.941, precision of 0.950, F1 of 0.931 and AUC of 0.965, followed by the late fusion and the early fusion, although the difference between the latter two is not statistically significant. Moreover, the results exhibit robustness, as indicated by the minimal standard deviations derived from 5 independent experiment runs.

Table 2: Model evaluations (the overall mean  $\pm$  standard deviation across 5 independent runs) for three fusion strategies (i.e., early fusion, late fusion and intermediate fusion) on the test set. Best results are highlighted in bold, and \* indicates that the comparison is statistically significant (p < 0.05).

_		(I ( 0 0 0 ) .		
	Fusion	Early Fusion	Late Fusion	Intermediate Fusion
	Strategies	(Fig. 3a)	(Fig. 3b)	(Proposed Fig. 4)
	Accuracy	$0.911 \pm 0.013$	$0.922 \pm 0.022$	$0.941\pm0.021^*$
	Precision	$0.918 \pm 0.008$	$0.920 \pm 0.011$	$0.950\pm0.010^*$
	F1	$0.901 \pm 0.009$	$0.901 \pm 0.014$	$0.931\pm0.015^*$
	AUC	$0.939 \pm 0.003$	$0.940 \pm 0.003$	$0.965\pm0.004^*$

## 3.3. Effects of Smoothing Levels

Figure 5 presents a comparative analysis of model performance across various smoothing levels of  $\beta$  in Eq. (3). At the outset, for  $\beta = 0$ , no label

smoothing yields a model accuracy of 0.921, precision of 0.915, F1 of 0.908 and AUC of 0.946. With an increase in  $\beta$ , the model's performance exhibits an initial ascent followed by a subsequent decline. When  $\beta=0.3$ , the model achieves the optimal performance across all metrics. Additionally, the results indicate that, with the implementation of label smoothing, for  $\beta \leq 0.5$ , the model consistently outperforms the one without label smoothing ( $\beta=0$ )..

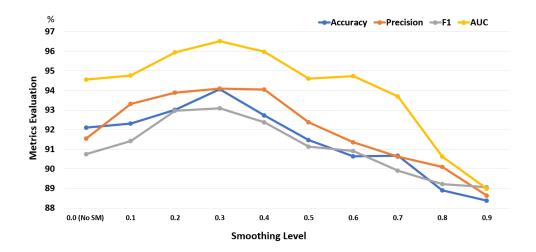


Figure 5: The evaluation of model performance (Accuracy, Precision, F1, and AUC scores) on the test set presented across varying smoothing levels ( $\beta$ ) ranging from 0 to 1 in increments of 0.1. Here,  $\beta = 0$  represents no smoothing.

# 3.4. Effects of Gated Attention Mechanism

Table 3 compares the performance of models with and without the gated attention mechanism. These models share an identical CNN architecture for feature extraction from each view but diverge in their feature fusion approaches, such as mean pooling and max pooling. The results demonstrate that the performance of the model with the gated attention mechanism is significantly better than of those without it. Additionally, statistical analysis indicates no significant difference between the performances of mean pooling and max pooling.

The gated attention mechanism also enhances interpretability, rendering the model's outputs more explainable to humans. Figure 6 shows examples with attention weights assigned by the proposed classifier for four glitch predictions. The magnitude of attention weights positively correlates with the contribution of each time view to the prediction. For instance, in Fig. 6 (a), (d) and (e), the model places greater emphasis on long time windows to predict glitches like Repeating Blips and No Glitch, while allocating higher attention to short time windows to predict glitches such as Blip and Koi Fish in Fig. 6 (b) and (c).

Table 3: Model evaluations (the overall mean  $\pm$  standard deviation across 5 independent runs) for classifiers with gated attention mechanism and without attention mechanism (using max pooling and mean pooling) on the test set. Best results are highlighted in bol, and \* indicates that the comparison is statistically significant (p < 0.05).

Attention	Proposed	No Attention	No Attention
Mechanism	(Gated Attention)	(Max Pooling )	(Mean Pooling)
Accuracy	$0.941\pm0.021^*$	$0.916 \pm 0.013$	$0.911 \pm 0.011$
Precision	$0.950\pm0.010^*$	$0.928 \pm 0.010$	$0.929 \pm 0.006$
F1	$0.931\pm0.015^*$	$0.908 \pm 0.011$	$0.903 \pm 0.008$
AUC	$0.965\pm0.004^*$	$0.948 \pm 0.002$	$0.949 \pm 0.001$

# 3.5. Comparison with Previous Classifiers

We compare the performance of the proposed model with the previously implemented classifier [20] in the O3 run on the same test set, as shown in Table 4. The proposed model shows superiority across all metrics, outperforming the previous model by 2.7% to 3.7%. Furthermore, a detailed examination of class-wise accuracies reveals that the proposed model provides more accurate predictions for the majority of glitch classes. In particular, the model excels in classes like Wandering Line, Repeating Blips, and Helix. While there are a few instances for which the previous classifier attains better outcomes, such as Koi Fish and Scattered Light, the differences remain relatively modest.

To obtain a deeper insight into this finding, we compare the t-SNE plots of these two classifiers within a two-dimensional space, as shown in Fig. 7. Compared to the previous classifier, it is evident that the feature distributions generated by the new classifier exhibit tighter clusters within the same glitch class and distinct separation among clusters of different classes. However, some degree of overlap persists between glitch classes with similar

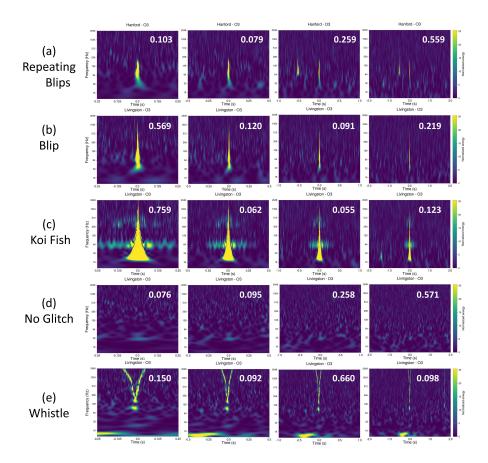


Figure 6: Illustrative examples of glitch classification with attention weights generated by the classifier for each time window view. Attention weights are shown in white in each subplot. A higher attention weight indicates that the model relies more on that particular view for predicting the glitch class.

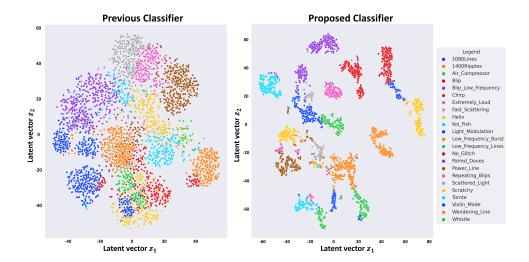


Figure 7: Comparison of t-SNE plots in a two-dimensional space (latent feature  $z_1$  and  $z_2$ ) between the previous classifier and the newly proposed classifier on the same test set. The proposed classifier (right) generates more tightly-clustered glitch features than the previous one (left) within the latent space.

morphologies, such as Extremely Loud and Koi Fish, Blip and Repeating Blips, Scattered Light and Low Frequency Lines.

Figure 8 presents the reliability plots of two classifiers applied to the test set. Both classifiers are overconfident in their predictions, evidenced by the majority of points residing below the diagonal line. Nevertheless, the proposed model demonstrates better calibration compared to the previous model, particularly in the high confidence range (> 0.8) – a crucial aspect for detector characterization and gravitational-wave data analysis. Additionally, the proposed model achieves a lower expected calibration error of 0.06, compared to the previous model's 0.09. This discrepancy indicates that the proposed model displays less bias and higher reliability than the previous classifier.

## 4. Discussion

In this study, we introduce a novel classifier for the Gravity Spy project, aiming to enhance the accuracy of glitch classification provided by the LIGO gravitational-wave detectors. Overall, our proposed classifier achieves accuracy of 0.941, precision of 0.950, F1 score of 0.931 and AUC of 0.965 on the

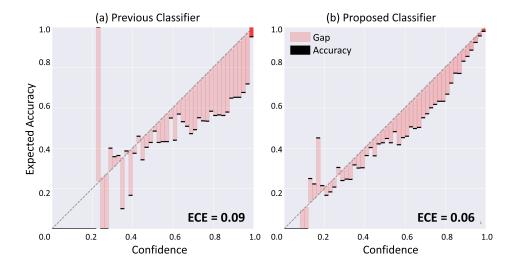


Figure 8: The reliability plots of the model output probabilities on the test set. The plots illustrate the observed fraction of positives against the predicted fraction of positives. The diagonal dotted line indicates perfect reliability. The expected calibration error (ECE) quantifies the average discrepancy between predicted probabilities and actual accuracy.

held-out set of glitches, outperforming prior classifiers integrated into the previous Gravity Spy workflow.

The proposed classifier exhibits superior glitch classification through various attributes. First, the use of the inception residual block empowers the model to deepen its architecture without overfitting concerns, thereby facilitating the extraction of more discriminative features unique to different glitch classes. The tightly-clustered distributions in Fig. 7 further substantiate this finding. Next, the intermediate fusion strategy outperforms early and late fusion due to its ability to capture and combine cross-view correlations, which aligns with findings from prior studies, such as using multiple time scales for anomaly detection [46] and multiple modalities for gesture recognition [47]. Moreover, label smoothing effectively tackles the issue of noisy labels within the training dataset by introducing a controlled level of uncertainty into label distributions during training, resulting in enhanced glitch predictions. However, it is crucial to choose the level of this uncertainty infusion appropriately, as excessive smoothing, illustrated in Fig. 5, can lead to label confusion within the model, hindering its ability to learn effectively.

Table 4: Comparison of class-wise classification accuracy and the overall performance (accuracy, precision, F1 and AUC scores) between the previous classifier and the proposed classifier. Better results are highlighted in bold, and \* indicates that the comparison is statistically significant (p < 0.05).

Class	Previous	Proposed	Class	Previous	Proposed
1080Lines	0.985	0.991	Paired Doves	0.849	0.882*
1400Ripples	0.967	0.979*	Power Line	0.987	0.994*
Air Compressor	0.921	0.954*	Repeating Blips	0.772	0.933*
Blip	0.980	0.991*	Scattered Light	0.929	0.924
Blip Low Frequency	0.926	0.928	Scratchy	0.862	0.957*
Chirp	0.913	0.926*	Tomte	0.963	0.984*
Extremely Loud	0.901	0.960*	Violin Mode	0.924	0.960*
Fast Scattering	0.884	0.890	Wandering Line	0.52	0.740*
Helix	0.732	0.891*	Whistle	0.957	0.972*
Koi Fish	0.921	0.917	Overall		
Light Modulation	0.918	0.947*	Accuracy	0.906	0.941*
Low Frequency Burst	0.962	0.981*	Precision	0.913	0.950*
Low Frequency Lines	0.818	0.916*	F1	0.904	0.931*
No Glitch	0.978	0.984	AUC	0.937	0.965*

The challenge of the black box issue in deep learning stems from the inherent lack of transparency in comprehending how these models arrive at decisions. The attention mechanism addresses this challenge by enabling the model to selectively focus on specific features of the input data, assigning varying levels of importance to different regions. By highlighting informative features and providing interpretable decisions, the attention mechanism enhances transparency in the decision-making process.

In our study, the model achieves interpretability by discerning which view contributes the most to the classification decision. Consequently, the attention mechanism significantly improves the overall model performance by selectively prioritizing informative aspects of each view and effectively capturing cross-view relationships. For example, as shown in Fig. 6, the model's extraction of more discriminative features from the 4 s and 2 s windows when classifying a glitch as Repeating Blips aligns with volunteers' experiences, reflecting their awareness that repeating patterns primarily manifest in longer time windows. In contrast, when identifying a single Blip class, the model focuses on extracting more specific features from the 0.5 s window, corresponding to the transient nature of this particular glitch category. Similarly, when dealing with cases resembling Blip or Koi Fish, the model

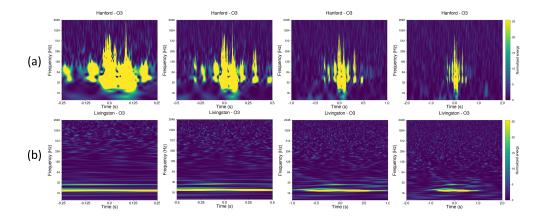


Figure 9: Examples of glitches with different predicted classes between the previous model and the proposed model. (a) Initially predicted as the Koi Fish class by the previous model, it is now classified as Extremely Loud by the proposed model. (b) Initially classified as the Scattered Light class by the previous model, it is now predicted as Low Frequency Lines by the proposed model.

directs greater attention to the 0.5 s window. Concerning No Glitch classification, the model's emphasis on the 4 s window is intuitive, aligning with the necessity to inspect the longest window time to confirm all background features or, in other words, the absence of glitches. It is essential to note that the emphasis on specific windows does not diminish the importance of those with low attention weights. Instead, the model effectively captures crucial information from all four windows, enabling it to make optimal glitch predictions. Potentially, the attention mechanism can guide volunteers and LIGO experts, enabling them to enhance their understanding of glitches and thereby facilitating a more effective characterization process.

Given its superior overall performance in comparison to the previous classifier, the proposed classifier has been integrated into the Gravity Spy system for the O4 run. This integration coincides with major infrastructure enhancements introduced in the O4 run [48, 22], aimed at achieving higher instrument sensitivity for the detection of gravitational-wave signals. However, this improvement in detector sensitivity is accompanied by a trade-off, i.e., a higher sensitivity to environmental and instrumental artifacts, which potentially leads to new glitch classes. In this context, the more discriminative features extracted by the proposed classifier, as shown in Fig. 7, hold the

potential to more effectively identify and categorize these new glitch classes.

An underperformance is observed in certain glitch classes, such as Koi Fish and Scattered Light as presented in Table 4. It is essential to note that this underperformance is not significant (p>0.05), approximately 0.5% compared to the previous model, which may potentially be attributed to random sampling variations in our test set. Besides, this underperformance is also influenced by the inherent complexity and variability within these classes. Given their prevalence and diverse morphologies, establishing strict boundaries for these classes poses a challenge, leading to expected variations in classifiers when distinguishing these intricacies. For instance, in Fig. 9, Koi Fish can be confused with Extremely Loud when the trigger is very loud [20] and Scattered Light can be confused with Low Frequency Lines as they share characteristics of low-frequency and long duration [49].

The previous Gravity Spy classifier has long grappled with the challenge of confidence overfitting. This issue can be attributed to various factors, including imbalances in our glitch classes, the presence of noisy labels in our training set, and the use of cross-entropy loss for updating the model parameters. These challenges pose the risk of hindering the model's ability to generalize effectively to new and unseen glitch data. To address this issue, our proposed classifier incorporates several techniques. We apply L2 regularization and incorporate smooth labeling into the cross-entropy loss to penalize extreme confidence. Moreover, we monitor accuracy for updating model parameters, extending beyond the sole reliance on cross-entropy loss, which tends to drive the predicted confidence close to 1.

To demonstrate the impact of confidence overfitting in the O4 run, we conduct a comparative analysis of the predicted confidence distributions between the previous model and the proposed model. This analysis encompasses a total of 2267 glitches extracted from May 24, 2023, to September 19, 2023, within the O4 run. As illustrated in Fig. 10 (a), the previous model exhibits persistent confidence overfitting, as most glitch classes receive predictions with confidence close to 1, except for Paired Doves and Light Modulation. In contrast, the proposed model, as shown in Fig. 10 (b), significantly mitigates this issue, providing predictions with varying confidence ranges across glitch classes. Specifically, for glitch classes like Extremely Loud and Low Frequency Lines, which exhibit less variation in morphologies within the O4 distribution compared to O3, the proposed model confidently expresses predictions with relatively high confidence. Conversely, for glitch classes such as Whistle, Wandering Line, Violin Mode, and Scratchy, the proposed model

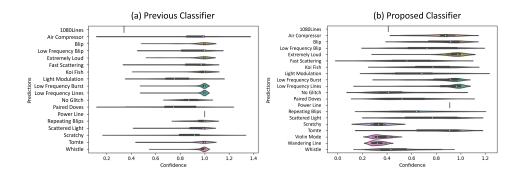


Figure 10: The density distribution plot of predicted confidence for O4 glitches between (a) the previous model and (b) the proposed model.

displays lower confidence in predictions, given the visually distinct morphologies in O4 compared to their distributions in O3. This distinction is crucial for LIGO experts during glitch analysis and characterization, allowing the exclusion of glitches with lower confidences to enhance the precision of their analyses. This, in turn, opens avenues to explore the potential discovery of new glitch classes within clusters exhibiting lower confidences. A good example is the identification of a new glitch class, named 589 Hz,  $^2$  depicted in Fig. 11, which occurred on May 30, 31, and June 1 of 2023 in the Hanford detector. The previous model predicted it as Whistle with an overall high confidence of  $0.976(\pm 0.059)$  on 267 collections, attributing this classification to its high-frequency nature similar to Whistle, while the proposed model predicted it with significantly lower confidence scores of  $0.463(\pm 0.138)$  for 195 cases classified as Whistle,  $0.428(\pm 0.137)$  for 49 cases classified as No Glitch and  $0.347(\pm 0.057)$  for 23 cases classified as Violin Mode, respectively.

This study has some limitations. First, as the classifier is trained under full supervision, its generalization capability to new classes is constrained. Consequently, if a novel glitch class arises during the O4 run, the classifier requires retraining on a dataset that incorporates the new glitch class. This process also involves the participation of volunteers who, with assistance from the Gravity Spy similarity search [16], identify and categorize these new glitch classes. Next, the training and test set used in this study are

<sup>&</sup>lt;sup>2</sup>https://www.zooniverse.org/projects/zooniverse/gravity-spy/talk/762/2986261

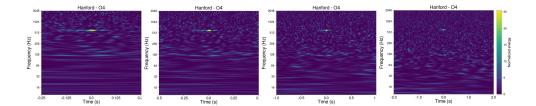


Figure 11: An example of a new glitch class, named 589 Hz in the O4 run. It is a narrow band line glitch near 589 Hz lasting less than 0.4 s, appearing in these spectrograms as bright yellow for 0.02 s to 0.2 s. It was classified as the Whistle class with a high confidence of 0.972 by the previous classifier and a low confidence of 0.348 by the proposed classifier.

still imperfect with noisy labels. To improve the reliability of the model, future efforts should be put into preparing a test set with validated ground truth, as confirmed by LIGO experts. Additionally, there is room for further optimization of the classifier for improved glitch classification. Future studies can enhance the model's training by incorporating additional information such as sensor locations, temporal attributes, and supplementary auxiliary data [15, 50, 51]. For instance, supplementary auxiliary channels, could assist in reducing potential Scattered Light, as illustrated in Fig. 9 (b), by aiding in the identification of ground motion [52]. Last, our model currently provides information on important views but does not specify which parts of the focused image trigger predictions. Our research will explore techniques like Grad-CAM [53] to precisely identify the influential regions within glitch images.

## 5. Conclusion

In this study, we present a novel classifier for the Gravity Spy project, specifically designed to enhance the accuracy of glitch classification within the LIGO gravitational-wave detectors. The proposed classifier demonstrates exceptional performance compared to the previous classifiers. Several factors contribute to the classifier's distinctiveness. Leveraging the inception residual block allows for deeper architecture exploration, facilitating the extraction of discriminative glitch features. The intermediate fusion strategy captures cross-view correlations more effectively. Label smoothing addresses noisy training labels, enhancing the generalization of glitch predictions. The

attention mechanism not only enhances model performance but also provides interpretability, potentially aiding volunteers and experts in understanding glitch classification decisions.

The new classifier has successfully been integrated into the ongoing O4 run, actively employed by both volunteers on the Zooniverse platform and LIGO experts. Its capacity to potentially help volunteers identify and categorize new glitch classes fits well with the larger context of improved detector sensitivity, where a more advanced model becomes essential for understanding the details of more intricate glitch patterns.

## Acknowledgements

The authors extend their gratitude to the community-science volunteers of the Gravity Spy project, whose dedicated efforts played a pivotal role in bringing this ambitious project to fruition. The authors also thank Man-Leong Chan for comments that improved the quality of this manuscript. 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 supported by National Science Foundation Grants PHY-0757058 and PHY-0823459. Gravity Spy is supported by NSF grants IIS-2106882, 2106896, 2107334, and 2106865. CPLB acknowledges support from Science and Technology Facilities Council (STFC) grant ST/V005634/1. ZD acknowledges support from the CIERA Board of Visitors Research Professorship and NSF grant PHY-2207945. This document has been assigned LIGO document number LIGO-P2300458.

# Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Data Availability Statement

The datasets from the Gravity Spy project, including machine learning and volunteer classification information, can be accessed through the Zenodo repositories [21, 30]. The corresponding codes are also accessible on GitHub through Gravity Spy.

#### References

- [1] A. Einstein, Approximative integration of the field equations of gravitation, Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.) 1916 (1) (1916).
- [2] T. L. S. Collaboration, J. Aasi, Advanced LIGO, Classical and Quantum Gravity 32 (7) (2015) 074001.
- [3] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, et al., Observation of gravitational waves from a binary black hole merger, Physical Review Letters 116 (6) (2016) 061102.
- [4] F. Acernese, M. Agathos, K. Agatsuma, et al., Advanced Virgo: a 2nd generation interferometric gravitational wave detector, Classical and Quantum Gravity 32 (2).
- [5] B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, et al., GW170814: A three-detector observation of gravitational waves from a binary black hole coalescence, Physical Review Letters 119 (2017) 141101.
- [6] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, et al., GWTC-1: a gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo during the first and second observing runs, Physical Review Letters 9 (3) (2019) 031040.
- [7] K. collaboration, KAGRA: 2.5 generation interferometric gravitational wave detector, Nature Astronomy 3 (1) (2019) 35–40.
- [8] L. S. Collaboration, V. Collaboration, K. Collaboration, R. Abbott, H. Abe, F. Acernese, K. Ackley, N. Adhikari, R. Adhikari, V. Adkins, et al., First joint observation by the underground gravitational-wave detector KAGRA with GEO 600, Progress of Theoretical and Experimental Physics 2022 (6) (2022) 063F01.
- [9] B. P. Abbott, R. Abbott, T. D. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. B. Adya, C. Affeldt, M. Agathos, et al., A guide

- to LIGO-Virgo detector noise and extraction of transient gravitational-wave signals, Classical and Quantum Gravity 37 (5) (2020) 055002.
- [10] B. P. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, M. Adamo, C. Adams, T. Adams, P. Addesso, et al., Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914, Classical and Quantum Gravity 33 (13) (2016) 134001.
- [11] D. Davis, J. S. Areeda, B. K. Berger, R. Bruntz, A. Effler, R. Essick, R. Fisher, P. Godwin, E. Goetz, A. Helmling-Cornell, et al., Ligo detector characterization in the second and third observing runs, Classical and Quantum Gravity 38 (13) (2021) 135014.
- [12] D. Davis, M. Walker, Detector characterization and mitigation of noise in ground-based gravitational-wave interferometers, Galaxies 10 (1) (2022) 12.
- [13] L. Nuttall, Characterizing transient noise in the LIGO detectors, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 376 (2120) (2018) 20170286.
- [14] M. Zevin, S. Coughlin, S. Bahaadini, E. Besler, N. Rohani, S. Allen, M. Cabero, K. Crowston, A. K. Katsaggelos, S. L. Larson, et al., Gravity Spy: Integrating Advanced LIGO detector characterization, machine learning, and citizen science, Classical and Quantum Gravity 34 (6) (2017) 064003.
- [15] M. Zevin, C. B. Jackson, Z. Doctor, Y. Wu, C. Østerlund, L. C. Johnson, C. P. Berry, K. Crowston, S. B. Coughlin, V. Kalogera, et al., Gravity Spy: Lessons learned and a path forward, arXiv preprint arXiv:2308.15530 (2023).
- [16] S. Coughlin, S. Bahaadini, N. Rohani, M. Zevin, O. Patane, M. Harandi, C. Jackson, V. Noroozi, S. Allen, J. Areeda, et al., Classifying the unknown: discovering novel gravitational-wave detector glitches using similarity learning, Physical Review D 99 (8) (2019) 082002.
- [17] L. Fortson, D. Wright, C. Lintott, L. Trouille, Optimizing the human-machine partnership with Zooniverse, arXiv preprint arXiv:1809.09738 (2018).

- [18] J. Dodge, G. Ilharco, R. Schwartz, A. Farhadi, H. Hajishirzi, N. Smith, Fine-tuning pretrained language models: Weight initializations, data orders, and early stopping, arXiv preprint arXiv:2002.06305 (2020).
- [19] S. Bahaadini, V. Noroozi, N. Rohani, S. Coughlin, M. Zevin, J. R. Smith, V. Kalogera, A. Katsaggelos, Machine learning for Gravity Spy: Glitch classification and dataset, Information Sciences 444 (2018) 172–186.
- [20] S. Soni, C. P. L. Berry, S. B. Coughlin, M. Harandi, C. B. Jackson, K. Crowston, C. Østerlund, O. Patane, A. K. Katsaggelos, L. Trouille, et al., Discovering features in gravitational-wave data through detector characterization, citizen science and machine learning, Classical and Quantum Gravity 38 (19) (2021) 195016.
- [21] J. Glanzer, S. Banagiri, S. Coughlin, S. Soni, M. Zevin, C. P. L. Berry, O. Patane, S. Bahaadini, N. Rohani, K. Crowston, et al., Data quality up to the third observing run of Advanced LIGO: Gravity Spy glitch classifications, Classical and Quantum Gravity 40 (6) (2023) 065004.
- [22] B. P. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, V. Adya, C. Affeldt, M. Agathos, et al., Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, Living reviews in relativity 23 (2020) 1–69.
- [23] S. Alvarez-Lopez, A. Liyanage, J. Ding, R. Ng, J. McIver, GSpyNet-Tree: A signal-vs-glitch classifier for gravitational-wave event candidates, arXiv preprint arXiv:2304.09977 (2023).
- [24] S. Jarov, S. Thiele, S. Soni, J. Ding, J. McIver, R. Ng, R. Hatoya, D. Davis, A new method to distinguish gravitational-wave signals from detector noise transients with Gravity Spy, arXiv preprint arXiv:2307.15867 (2023).
- [25] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv preprint arXiv:1409.1556 (2014).
- [26] S. C. Huang, A. Pareek, S. Seyyedi, I. Banerjee, M. P. Lungren, Fusion of medical imaging and electronic health records using deep learning: a systematic review and implementation guidelines, NPJ digital medicine 3 (1) (2020) 136.

- [27] F. Robinet, N. Arnaud, N. Leroy, A. Lundgren, D. Macleod, J. McIver, Omicron: a tool to characterize transient noise in gravitational-wave detectors, SoftwareX 12 (2020) 100620.
- [28] S. Chatterji, L. Blackburn, G. Martin, E. Katsavounidis, Multiresolution techniques for the detection of gravitational-wave bursts, Classical and Quantum Gravity 21 (20) (2004) S1809.
- [29] S. K. Chatterji, The search for gravitational wave bursts in data from the second LIGO science run, Ph.D. thesis, Massachusetts Institute of Technology (2005).
- [30] M. Zevin, S. Coughlin, E. Chase, S. Allen, S. Bahaadini, C. Berry, K. Crowston, M. Harandi, C. Jackson, V. Kalogera, et al., Gravity spy volunteer classifications of LIGO glitches from observing runs O1, O2, O3a, and O3b (2022).
- [31] S. Bahaadini, N. Rohani, S. Coughlin, M. Zevin, V. Kalogera, A. K. Katsaggelos, Deep multi-view models for glitch classification, in: 2017 ieee international conference on acoustics, speech and signal processing (ICASSP), IEEE, 2017, pp. 2931–2935.
- [32] C. Szegedy, S. Ioffe, V. Vanhoucke, A. Alemi, Inception-v4, inception-resnet and the impact of residual connections on learning, in: Proceedings of the AAAI conference on artificial intelligence, Vol. 31, 2017.
- [33] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, A. Rabinovich, Going deeper with convolutions, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2015, pp. 1–9.
- [34] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [35] M. Ilse, J. Tomczak, M. Welling, Attention-based deep multiple instance learning, in: International conference on machine learning, PMLR, 2018, pp. 2127–2136.

- [36] Y. Wu, A. Schmidt, E. Hernández-Sánchez, R. Molina, A. K. Katsaggelos, Combining attention-based multiple instance learning and gaussian processes for CT hemorrhage detection, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 582–591.
- [37] Y. Wu, F. M. Castro-Macías, P. Morales-Álvarez, R. Molina, A. K. Katsaggelos, Smooth attention for deep multiple instance learning: Application to ct intracranial hemorrhage detection, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2023, pp. 327–337.
- [38] M. López-Pérez, A. Schmidt, Y. Wu, R. Molina, A. K. Katsaggelos, Deep gaussian processes for multiple instance learning: Application to CT intracranial hemorrhage detection, Computer Methods And Programs In Biomedicine 219 (2022) 106783.
- [39] R. Müller, S. Kornblith, G. E. Hinton, When does label smoothing help?, Advances in neural information processing systems 32 (2019).
- [40] C. Cortes, M. Mohri, A. Rostamizadeh, L2 regularization for learning kernels, arXiv preprint arXiv:1205.2653 (2012).
- [41] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, The journal of machine learning research 15 (1) (2014) 1929–1958.
- [42] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- [43] T. K. Kim, T test as a parametric statistic, Korean journal of anesthesiology 68 (6) (2015) 540–546.
- [44] J. Bröcker, L. A. Smith, Increasing the reliability of reliability diagrams, Weather and forecasting 22 (3) (2007) 651–661.
- [45] C. Guo, G. Pleiss, Y. Sun, K. Q. Weinberger, On calibration of modern neural networks, in: International conference on machine learning, PMLR, 2017, pp. 1321–1330.

- [46] W. Wang, F. Chang, H. Mi, Intermediate fused network with multiple timescales for anomaly detection, Neurocomputing 433 (2021) 37–49.
- [47] A. Roitberg, T. Pollert, M. Haurilet, M. Martin, R. Stiefelhagen, Analysis of deep fusion strategies for multi-modal gesture recognition, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2019, pp. 0–0.
- [48] C. Cahillane, G. Mansell, Review of the Advanced LIGO gravitational wave observatories leading to observing run four, Galaxies 10 (1) (2022) 36.
- [49] T. Accadia, F. Acernese, F. Antonucci, P. Astone, G. Ballardin, F. Barone, M. Barsuglia, T. S. Bauer, M. Beker, A. Belletoile, et al., Noise from scattered light in Virgo's second science run data, Classical and Quantum Gravity 27 (19) (2010) 194011.
- [50] P. Nguyen, R. Schofield, A. Effler, C. Austin, V. Adya, M. Ball, S. Banagiri, K. Banowetz, C. Billman, C. Blair, et al., Environmental noise in Advanced LIGO detectors, Classical and Quantum Gravity 38 (14) (2021) 145001.
- [51] R. E. Colgan, K. R. Corley, Y. Lau, I. Bartos, J. N. Wright, Z. Márka, S. Márka, Efficient gravitational-wave glitch identification from environmental data through machine learning, Physical Review D 101 (10) (2020) 102003.
- [52] S. Soni, C. Austin, A. Effler, R. Schofield, G. González, V. Frolov, J. Driggers, A. Pele, A. Urban, G. Valdes, et al., Reducing scattered light in LIGO's third observing run, Classical and Quantum Gravity 38 (2) (2020) 025016.
- [53] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, Grad-CAM: Visual explanations from deep networks via gradient-based localization, in: 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 618–626.