

One flexible model for multiclass gravitational wave signal and glitch generation

Tom Dooney¹, R. Lyana Curier¹, Daniel Stanley Tan¹, Melissa Lopez^{2,3},
Chris Van Den Broeck^{2,3} and Stefano Bromuri¹

¹Faculty of Science, Open Universiteit, Valkenburgerweg 177, 6419 AT Heerlen, The Netherlands

²Institute for Gravitational and Subatomic Physics (GRASP), Utrecht University,
Princetonplein 1, 3584 CC Utrecht, The Netherlands

³Nikhef, Science Park 105, 1098 XG, Amsterdam, The Netherlands



(Received 7 February 2024; accepted 4 June 2024; published 16 July 2024)

Simulating realistic time-domain observations of gravitational waves (GWs) and other events of interest in GW detectors, such as transient noise bursts called glitches, can help in advancing GW data analysis. Simulated data can be used in downstream data analysis tasks by augmenting datasets for signal searches, balancing datasets for machine learning applications, validating detection schemes, and constructing mock data challenges. In this work, we present a conditional derivative GAN (cDVGAN), a novel conditional model in the generative adversarial network framework for simulating multiple classes of time-domain observations that represent gravitational waves (GWs) and detector glitches. cDVGAN can also generate generalized hybrid samples that span the variation between classes through class interpolation in the conditioned class vector. cDVGAN introduces an additional player into the typical 2-player adversarial game of GANs, where an auxiliary discriminator analyzes the first-order derivative time series. Our results show that this provides synthetic data that better capture the features of the original data. cDVGAN conditions on three classes in the time domain, two denoised from LIGO *blip* and *tomte* glitch events from its third observing run (O3), and the third representing binary black hole (BBH) mergers. Our proposed cDVGAN outperforms four different baseline GAN models in replicating the features of the three classes. Specifically, our experiments show that training convolutional neural networks (CNNs) with our cDVGAN-generated data improves the detection of samples embedded in detector noise beyond the synthetic data from other state-of-the-art GAN models. Our best synthetic dataset yields as much as a 4.2% increase in *area-under-the-curve* (AUC) performance, maintaining the same CNN architecture, compared to synthetic datasets from baseline GANs. Moreover, training the CNN with class-interpolated hybrid samples from our cDVGAN outperforms CNNs trained only on the standard classes, when identifying real samples embedded in LIGO detector background between signal-to-noise ratios ranging from 1 to 16 (4% AUC improvement for cDVGAN). We also illustrate an application of cDVGAN in a data augmentation example, showing that it is competitive with a traditional augmentation approach. Lastly, we test cDVGAN's BBH signals in a fitting-factor study, showing that the synthetic signals are generally consistent with the semianalytical model used to generate the training signals and the corresponding parameter space.

DOI: [10.1103/PhysRevD.110.022004](https://doi.org/10.1103/PhysRevD.110.022004)

I. INTRODUCTION

The first detection of a gravitational wave (GW) from a binary black hole (BBH) merger in 2015 ushered in a new era of astronomy and cosmology [1]. Since then, over three observing runs (O1, O2, O3), advanced LIGO [2], and Virgo [3] detectors have made confident detections of 90 compact binary coalescence (CBC) events, as reported in the gravitational wave transient catalogs GWTC-1, GWTC-2, GWTC-3 [4–6]. With the introduction of KAGRA [7], Japan's underground detector, towards the end of O3 and the O4 run currently underway, hundreds of more detections are expected from the enhanced sensitivity of GW detectors [8].

Ongoing upgrades to advanced detector systems will give rise to new challenges in gravitational wave (GW) data analysis, particularly with the introduction of next-generation GW detectors, such as the Einstein Telescope (ET) [9] and Cosmic Explorer [10]. GW detection rates from all sources are expected to significantly increase; i.e., it is estimated the ET will detect on the order of $8 \times 10^4 \text{ y}^{-1}$ BBH mergers [11] and $7 \times 10^4 \text{ y}^{-1}$ binary neutron star (BNS) mergers [12]. This could lead to over 400 compact binary coalescence (CBC) events daily. The enhanced detectors are also expected to detect new astrophysical sources of GWs [13,14].

Aside from genuine astrophysical signals, the heightened sensitivity of GW detectors is expected to exacerbate issues

relating to “glitches” [8,15–17]. Glitches are non-Gaussian transient noise artifacts that can resemble astrophysical signals, hindering GW data analysis and increasing false positives [18–27]. Glitches are unmodeled noise events, stemming from environmental or instrumental factors, with some sources remaining unidentified [28]. Unlike modeled CBC events, which are detected through matched-filtering [29–31], glitches are detected by model-free detection algorithms that scrutinize excess power in the time-frequency (spectrogram) representation to distinguish them from the detector background [32].

To meet the challenges of advanced detector systems, machine learning algorithms have become increasingly popular in the GW physics community [33–36]. In the case of glitches, studies have largely focused on spectrogram representations for identification due to their unmodeled nature. For example, *Gravity Spy* [37,38] have made significant strides in characterizing spectrogram representations of glitches by combining machine learning and citizen science. Multiple studies have employed machine learning extensively to improve classification accuracy on *Gravity Spy* spectrograms [39–41]. Others have leveraged the generative adversarial network (GAN) [42] framework for generating *Gravity Spy* spectrograms to augment glitch datasets for further improvement [43,44]. However, relying on computationally expensive spectrogram transformations to identify and simulate unmodeled events like glitches may not always be feasible as GW detector technology improves.

In this study, we develop a generative modeling framework for diverse classes of time-domain observations in GW detectors. Generating time-domain representations of GWs and glitches offers various advantages, such as their flexibility for experimental purposes and low dimensionality requiring less computational expense. Simulated data can be used in downstream applications such as data augmentation and class balancing for machine learning applications, validating detection schemes via software injections [45–47], and constructing mock data challenges. Furthermore, where unmodeled transients like glitches can be isolated from the background, transforming time-series to spectrograms is straightforward, while the reverse is challenging due to background noise captured in spectrograms.

An approach to simulating time-domain glitch events is implemented using the *gengli* glitch generator [48]. The authors implement a Wasserstein GAN (WGAN) [49] on blip glitches extracted from detector backgrounds using *BayesWave* [50]. They show that it is possible to isolate blips from their surroundings and learn their underlying distribution in the time-domain with WGANs. Another study [51] implements a conditional GAN (cGAN) called McGANn to simulate five waveform classes analogous to GW bursts. Aside from generating five distinct classes, their GAN can generate class-interpolated, hybrid samples.

In this work, we propose a novel conditional derivative GAN (cDVGAN) that simulates different classes of

time-domain observations of LIGO glitch classes and/or astrophysical waveforms. We condition cDVGAN on three classes; two *Gravity Spy* glitch classes called *blip* and *tomte*, and a third represented by BBH signals. Aside from learning realistic distributions of diverse classes in one flexible model, cDVGAN can generate hybrid samples that traverse the variation between the learned classes by manipulating the user-controlled class vector. After the training phase, the models can generate hundreds of thousands of in-class or hybrid samples in a matter of seconds, per the user’s preference. cDVGAN is intended for use in a next-generation glitch generator, once investigations into *BayesWave* or other denoising algorithms can provide reliable time-domain representations of other glitch types.

To show the utility of our generated (synthetic) data, we implement an experiment that uses it for training a convolutional neural network (CNN) detection algorithm to identify real data from the GAN training data distribution in additive LIGO detector noise. Our results indicate that cDVGAN can better capture the features of the data than other state-of-the-art GAN models by incorporating adversarial feedback on the first-order derivatives using an auxiliary discriminator. Furthermore, we show that GAN-generated hybrid samples can be useful for training detection algorithms to identify real data in noise beyond the standard GW and glitch classes.

Since cDVGAN is primarily intended for use as a glitch generator, the astrophysical nature of the BBH signals is considered arbitrary for most of the experiments. However, we also confirm our cDVGAN data against a semianalytical model used in GW searches, exploring cDVGAN’s application as a BBH signal generator. To this end, we implement a fitting-factor study to evaluate the faithfulness of our generated data against the templates of a template bank. The results show that cDVGAN’s synthetic signals are generally consistent with signals from the waveform routine used to generate the original training signals.

This paper is structured as follows. In Sec. II, we discuss concepts relating to GANs. In Sec. III, we present the cDVGAN architecture and training schemes, the datasets and preprocessing, and experimental details. In Sec. IV, we present the experimental results involving a simple search for real GWs and glitches in additive detector background using synthetic data from five different GAN models. Finally, Sec. V discusses the conclusions of this research.

II. GENERATIVE ADVERSARIAL NETWORKS

A. Wasserstein GANs

GANs are a class of machine learning algorithms that generate realistic synthetic data. They consist of two neural networks: a discriminator (also known as a critic) that distinguishes between real and synthetic data, and a generator that generates synthetic data that can fool the discriminator. GAN approaches suffer from stability issues

such as vanishing gradients stimulating numerous studies centered around methods to stabilize the training process. Wasserstein GAN [49] is a particular variant that addresses these issues. It uses the Wasserstein-1 distance ($W1$) as the loss function to measure the similarity between the real and synthetic distributions. $W1$ is fully differentiable and increases monotonically while never saturating, removing the issue of vanishing gradients. Under this paradigm, the optimization problem can be formulated as

$$\theta_{\text{opt}} = \arg \min_{\theta} \max_{\phi: \|D(x, \phi)\|_1 \leq 1} L(\phi, \theta), \quad (1)$$

where the maximum is taken over all 1-Lipschitz functions D and with L defined as

$$L(\phi, \theta) = \mathbb{E}_{x \sim P_x} [D(x, \phi)] - \mathbb{E}_{\hat{x} \sim P_{\hat{x}}} [D(\hat{x}, \phi)], \quad (2)$$

where $\hat{x} = G(z, \theta)$ and z is a batch of the generator's latent vector. D and G refer to the discriminator and the generator with parameters ϕ and θ , respectively. $\mathbb{E}_{x \sim P_x}$ averages over a batch of real samples x from the real distribution P_x , while $\mathbb{E}_{\hat{x} \sim P_{\hat{x}}}$ averages over a batch of generated samples \hat{x} from the synthetic distribution $P_{\hat{x}}$. Equation (1) requires a constraint of 1-Lipschitz continuity on D [49]. This can be accomplished by adding a regularization penalty called the gradient penalty (GP) to the discriminator loss [52–54]. The discriminator loss then becomes

$$L_D = -L(\phi, \theta) + \lambda GP(\phi), \quad (3)$$

with

$$GP(\phi) = \mathbb{E}_{\hat{x} \sim P_{\hat{x}}} [(\|\nabla_x D(\hat{x}, \phi)\|_2 - 1)^2], \quad (4)$$

and where λ represents the regularization hyperparameter, $\|\cdot\|_2$ represents the L^2 -norm and \hat{x} is a randomly sampled point between the real and synthetic data. When updating the generator, errors are propagated through the entire network, from D to G . Naturally, for the generator updates are made only on generated samples from G . The generator loss is written as

$$L_G(\phi, \theta) = -\mathbb{E}_{\hat{x} \sim P_{\hat{x}}} [D(\hat{x}, \phi)]. \quad (5)$$

B. Conditional GANs

Conditional GANs (cGANs) [55] allow finer control over the generated data by providing extra information to both the generator and discriminator. For example, we can specify the class of the generated data by providing class label, c . The training data and class labels are taken from a joint distribution $P_{\text{data}}(x, c)$. When generating synthetic samples, the class vector \hat{c} and the GAN latent vector distribution P_z , are sampled independently.

In the original cGAN paper [55], conditional information was provided to the network by naively concatenating the class information to the first layers of the GAN components. In the case of time series, this corresponds to concatenating it to the z vector before passing it to the generator, and the input sample before passing it to the discriminator.

Miyato and Koyama [56] developed a more effective conditioning method using a projection in the discriminator between the conditional features and the features extracted from the input before the discriminator output. The projection output measures the similarity between the condition and the discriminator's feature vector and is added to the discriminator output. The output of the discriminator is modified from

$$D(x, c) = g(h(x, c)) \quad (6)$$

to

$$D(x, c) = c^T V h(x) + g(h(x)), \quad (7)$$

where V is an embedding matrix, h is the discriminator's feature vector and g is the discriminator output layer (see Fig. 2). This approach has been shown to improve the quality of class conditional generation using GANs. Note that the generator is still conditioned by concatenating the class-embedded vector to the latent input, similar to the original cGAN paper.

III. METHODS

A. Conditional derivative GAN (cDVGAN)

The adversarial training process for GANs is known to be volatile. Models can suffer mode collapse, where only a few realistic samples are learned, while they can also fail to converge at all. These issues often occur when one model component (generator/discriminator) begins to dominate the other during the training phase. These problems are exacerbated when conditioning on multimodal distributions of GW and glitch time series. In this section, we introduce two new cGAN designs, cDVGAN and cDVGAN2, that can help overcome these limitations.

In cDVGAN, two discriminators are applied to two different representations of the data instead of the usual single discriminator. The first discriminator is applied to the original samples as in a conventional GAN, while the second discriminator is applied to the corresponding derivative samples. A high-level diagram of cDVGAN can be seen in Fig. 1. We have shown in our previous work that the derivative discriminator leads to increased training stability of the model components and minimizes high-frequency artefacts in the GAN output, generating smoother and more faithful data [57]. Under this scheme, the generator loss is calculated as a linear combination of the two discriminator outputs. Equation (5) can be rewritten as

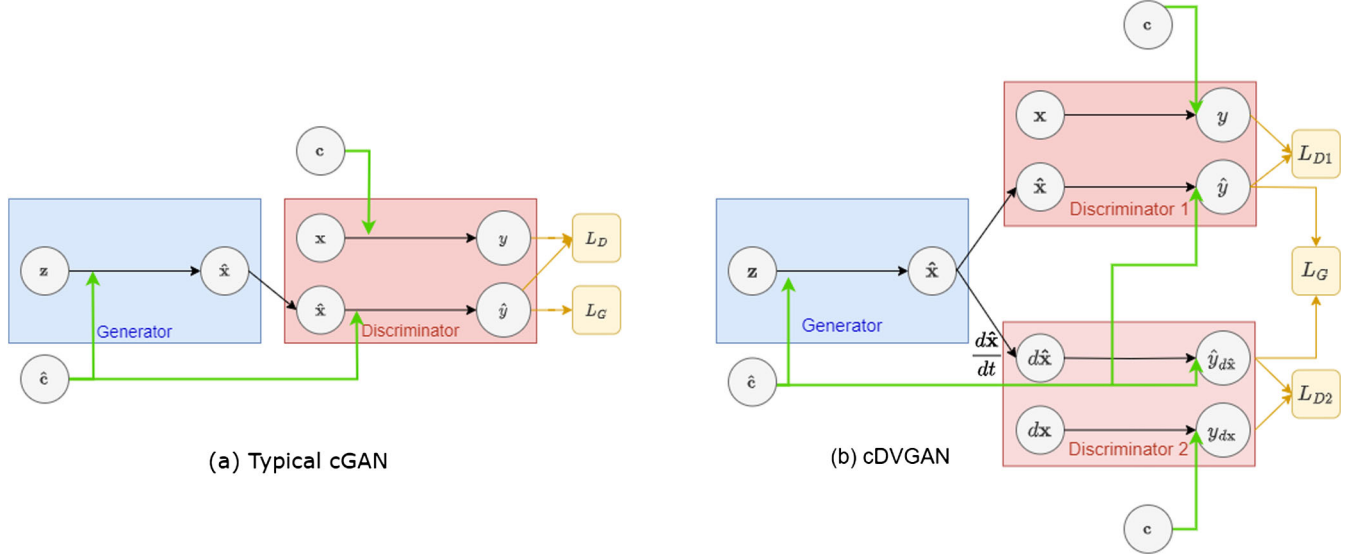


FIG. 1. Diagrams of a typical cGAN architecture (left), comprising one discriminator, and cDVGAN (right), comprising two discriminators. Class vectors c (real) and \hat{c} (fake), are fed to all model components in both cases. An intermediate derivative calculation is observed in the cDVGAN plot, where the derivative of the synthetic sample is calculated. cDVGAN2 includes yet another discriminator applied to second-order derivatives. In cDVGAN and cDVGAN2, the total generator loss is calculated as a linear combination of the discriminator losses applied to synthetic samples.

$$L_G(\phi_1, \phi_2, \theta) = -\eta_1 \mathbb{E}_{\hat{x}_1 \sim P_{\hat{x}_1}} [D_1(\hat{x}_1, \phi_1, \hat{c})] - \eta_2 \mathbb{E}_{\hat{x}_2 \sim P_{\hat{x}_2}} [D_2(\hat{x}_2, \phi_2, \hat{c})], \quad (8)$$

where D_1 and D_2 , represent the first and second discriminator, respectively. Here, \hat{x}_1 represents synthetic samples while $\hat{x}_2 = d\hat{x}_1/dt$ represents the corresponding derivatives and \hat{c} represents the class vector for the synthetic samples. $P_{\hat{x}_1}$ and $P_{\hat{x}_2}$ are the distributions of the two representations of generated data, and η_1 and η_2 are hyperparameters that

control the relative strength of the discriminator losses. The cDVGAN method is not restricted to two discriminators. We extend cDVGAN to cDVGAN2, which includes a third discriminator that is applied to second-order derivatives. The models are identical except for the addition of the second-order derivative discriminator in cDVGAN2, while the second-order derivative discriminator is identical to the first-order discriminator except for the input size. Additional representations of the data (e.g., time-frequency representations) can also be provided to additional discriminators, depending on the problem. For k discriminators applied to k representations of the data, the generator loss is written generally as

$$L_G(\phi_1, \dots, \phi_k, \theta) = -\sum_{i=1}^k \eta_i \mathbb{E}_{\hat{x}_i \sim P_{\hat{x}_i}} [D_i(\hat{x}_i, \phi_i, \hat{c})]. \quad (9)$$

Our cDVGAN and cDVGAN2 models are conditioned via projection, as described in Sec. II B. Since hyperparameter optimization is not the focus of this research, the hyperparameters $\eta_1 = \eta_2 = 0.5$ in cDVGAN and $\eta_1 = \eta_2 = \eta_3 = 0.33$ in cDVGAN2, meaning all discriminators contribute equally to the respective generator losses. The discriminators are updated 5 times for each generator update. The full model architecture can be viewed in the Appendix.¹

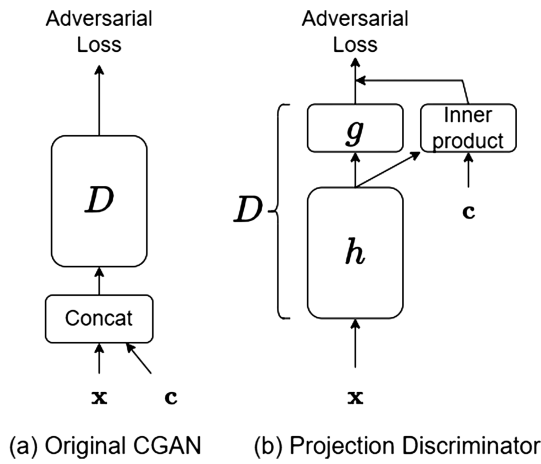


FIG. 2. A comparison of the discriminators from the original cGAN paper (used for McGANn and McDVGANn) and the projection discriminator (used for cWGAN, cDVGAN, cDVGAN2).

¹PYTHON code can be found at https://git.ligo.org/tom.dooney/cdvgan_paper.




FIG. 3. Examples of blip (top), tomte (middle), and BBH signals (bottom) used to train GAN models.

B. Training data and preprocessing

The GAN models in this study are conditioned on three different classes: two classes that are derived from *Gravity Spy* glitch classes called blip and tomte, and a third BBH signal class. The astrophysical nature that the BBH signals represent is mainly ignored during experimentation, and they are generally treated as another class of transient, “glitchlike” time series for experimental purposes. Examples of the three classes can be observed in Fig. 3. Blip glitches have a characteristic time-frequency morphology of a symmetric teardrop shape in the range [30, 500] Hz with short durations (~ 0.04 s). They appear in LIGO Livingston and LIGO Hanford, Virgo and GEO 600 [58]. Due to their abundance and form, they hinder both the unmodeled burst and modeled CBC searches, with particular emphasis on compact binaries with large total mass, highly asymmetric component masses, and spins antialigned with the orbital angular momentum [18,27]. Tomte glitches are also short duration (~ 0.25 s) with characteristic triangular morphology. Since both blip and tomte glitches have no clear correlation to the auxiliary channels, they cannot be removed from astrophysical searches. The BBH class represents the inspiral and merger of a binary black hole system. All samples are of length 1,024 and have a sampling rate of 4,096 Hz, corresponding to 0.25s of data. The GANs are trained on 7,500 samples (2,500 samples from each class).

The blip and tomte datasets are constructed using confidences from *Gravity Spy* applied to the glitch triggers in LIGO’s third observing run (O3) [59]. Only blip and tomte events with *Gravity Spy* confidences of $c_{\text{GS}}^1 \geq 0.9$ for their respective class are used in this study and are extracted using *GWpy*’s [60] `fetch_open_data` method, which provides an interface to the *GWOSC*² [61] data archive. The glitches are surrounded by stationary and uncorrelated noise, which would hinder the learning of GAN models.

²The Gravitational Wave Open Science Centre.

To avoid the computational expense of *BayesWave*, as done in [48], the glitches are isolated from the background using Savitsky-Golay [62] filters. This requires the following preprocessing steps:

- (1) Firstly, 20s of strain data is extracted, centred around the glitch GPS time provided by *Gravity Spy*.
- (2) The data are whitened, and a bandpass filter between (20, 350) Hz is applied to the 20s of data.
- (3) The data are cropped at 8,192 data points, centred around the GPS time of the glitch, corresponding to 2s of data.
- (4) A window of 100 data points is isolated around the centre of the glitch, and two consecutive Savitsky-Golay filters with a polynomial of order 3 are applied to each side around the glitch center iteratively, with window sizes of 501 and 301, respectively.
- (5) The two smoothed sides and the unsmoothed central peak are then concatenated, followed by applying three additional Savitsky-Golay filters iteratively with window sizes 41, 31, and 21 to the entire concatenated sample. The entire sample is finally cropped at 1,024 data points around the event GPS time, centered around 0 and rescaled to $(-1, 1)$.

The results of the above preprocessing steps are shown in Fig. 4. The extensive use of Savitsky-Golay filters ensures that high-frequency noise artifacts are removed while preserving the overall shape. An analysis was made to investigate whether the characteristics of blips and tomtes are preserved after filtering (Appendix). *Gravity Spy* generally classifies the preprocessed samples as their correct class when embedded in detector background for the signal-to-noise ratio (SNR) ranges considered in experiments (see Sec. III C), indicating that their morphologies are analogous to glitches.

All BBH signals are simulated with *PyCBC* [63] using the *IMRPhenomD* waveform routine from *LALSuite* [64], which generates the inspiral, merger, and ringdown of a BBH waveform. The component masses are restricted to the range of $[30, 160]M_{\odot}$ with a spin of zero and fixing $m_1 > m_2$ and using only plus polarization.

We show the diversity of the three classes in Fig. 5, where the dataset is represented by three separable clusters in a reduced principal component space.

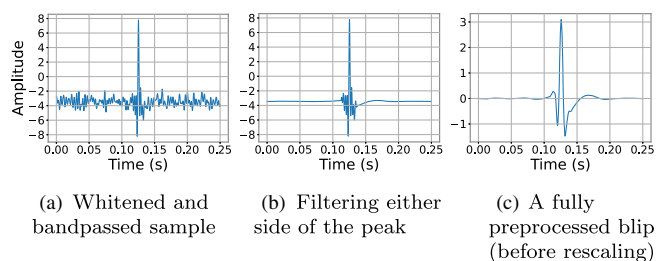


FIG. 4. Visualizations of the preprocessing steps applied to a blip glitch event.




FIG. 5. A plot of the first three principal components of the original samples. The separability of the three classes in this compressed representation indicates the diversity of the data.

C. Experimental procedure

1. GAN benchmarks

This work uses ablation studies to compare cDVGAN and its cDVGAN2 extension with three other baseline GAN models in their ability to generate data useful for training detection algorithms (see Sec. III C 3). An ablation study in machine learning is a systematic experimentation technique used to understand the contribution of individual components of a model to its overall performance. In this case, it involves selectively disabling auxiliary first and second-order derivative discriminators during the training of GAN models to investigate if they are successful in improving the features captured in synthetic GW signals and glitches. To construct an appropriate ablation study, the first benchmark is a conditional variant of a vanilla Wasserstein GAN (cWGAN), which comprises the same architecture as cDVGAN except for the derivative discriminator. This allows us to investigate the effect of including a first-order derivative discriminator during GAN training. cDVGAN2 allows us to investigate the effect of including a second-order derivative discriminator on top of the first-order derivative discriminator.

The second benchmark is McGANn [51] since it successfully replicates the features of simulated GW bursts. As a third benchmark, we developed a modified McGANn model, called McDVGANn, that uses a second auxiliary discriminator applied to first-order derivatives, similar to cDVGAN. The architecture is identical to McGANn except for the addition of a derivative discriminator. In McDVGANn, the derivative discriminator is identical to the base discriminator except for the input size. This benchmark investigates whether the idea of derivative discriminators can generalize

to other GAN models. For a more detailed description of the baselines, see Appendix.

McGANn and McDVGANn are conditioned using concatenation similar to their paper [51] and are trained with an Adam optimizer, binary cross-entropy loss function, and learning rate of 2×10^{-4} . Conversely, cWGAN and cDVGAN2 are conditioned via projection similarly to cDVGAN and are trained with RMSProp, a Wasserstein loss function and a learning rate of 2×10^{-4} . All models are trained for 500 epochs with a batch size of 512. Unless otherwise specified, all GANs are trained using the same standard hyperparameters as prior works.

2. GAN-generated datasets

We construct three different datasets from each GAN by sampling the respective generator’s class space, as suggested by [51]. The three variants of GAN-generated datasets are as follows:

- (1) *Vertex*: The vertex class space corresponds to the vertices of the three-dimensional class space. The locations of the vertex class space are the same as the training set class space locations used to train the GAN models and are the closest representation to the training set. The vertex class vector is one-hot encoded corresponding to one of the three GAN training classes, e.g., $[1, 0, 0]$ corresponds to the blip class.
- (2) *Simplex*: The simplex class space corresponds to points on a $k = 2$ simplex (2D triangle) in the case of three classes. Class vectors are constructed by sampling points uniformly on this simplex. The simplex can be considered the simplest surface that intersects all three training classes, and all simplex class vectors sum to 1. Variations are observed in the samples, with some having characteristics that strongly resemble the training classes, due to one class dominating the others. The simplex dataset is a superset of the vertex dataset and represents synthetic data outside the training data distribution.
- (3) *Uniform*: For the uniform dataset, each entry in the class vector is uniformly sampled from $U[0, 1]$, corresponding to sampling uniformly within a cube with dimensions $1 \times 1 \times 1$. The uniform dataset is a superset of the simplex and vertex datasets and, like the simplex dataset, represents data from outside of the training data distribution. The uniform dataset exhibits the largest variety since it explores regions of the class space further than the simplex dataset relative to the training set vertices.

Figure 6 shows examples from each of the cDVGAN-generated datasets (also in Appendix). Figure 7 shows plots of the first three principal components (PCs) of real and GAN-generated samples from cDVGAN. The vertex samples [Fig. 7(a)] generally match their corresponding classes from the real data distribution. Figure 7(b) shows that the





FIG. 6. Examples of vertex (top), simplex (middle) and uniform samples (bottom) from cDVGAN. The corresponding class vector is shown above each sample.

simplex and uniform hybrid datasets populate spaces between the class clusters, while Fig. 7(c) shows that the uniform dataset covers a larger part of the PC space than the simplex dataset. This is intuitive since the simplex space is a subset of the uniform space.

3. Downstream search with CNNs

The experiments followed in this study investigate the effectiveness of GAN-generated data for training CNNs to detect real data in additive Gaussian noise. CNNs are a class of deep learning model commonly used in computer vision. However, they can also be applied to other spatially adjacent data types like time series [65] and have been applied to GW detector strain data to detect merging black holes [34].

The CNNs’ objective is to perform the binary classification of two classes: samples in additive detector noise and detector noise only, examples of which are shown in



(a) Pure LIGO detector noise (0 class) (b) Detector noise + glitch (1 class)

FIG. 8. Examples of the two classes predicted by CNN models. The injected sample on the right is scaled to an SNR of 8 before injection (shown in orange for clarity).

Fig. 8. It takes a time-series input of dimension 1,024, representing 0.25 s of LIGO strain data. The CNN architecture is kept constant for training with each dataset and can be viewed in the Appendix.

Better synthetic data will result in CNNs with better detection efficiency on a real data subset from the GAN training data distribution. We also investigate if training CNNs with GAN-generated hybrid datasets (simplex and uniform) can improve the detection efficiency beyond training solely with the standard three classes.

The training and testing samples are injected additively into detector noise from Hanford (H1) or Livingston detectors (L1) during O3 for each of the above datasets. The detector noise is extracted from *GWOSC*, using trigger times for *Gravity Spy’s No_glitch* class, extracting 14s around each *No_glitch* GPS time. Although it is not investigated whether other glitches are present in the noise, using *Gravity Spy’s No_glitch* class should guarantee that most of the samples contain stationary noise after whitening.

The background is sampled at 4,096 Hz, similar to the GAN training data. The LIGO detector noise is first whitened using *PyCBC’s* whitening function, with the first and last 2.5 s removed due to artifacts at the noise




FIG. 7. The first three principal components (PCs) of real and GAN-generated samples from cDVGAN. The vertex samples from cDVGAN generally match the real samples in the PC space while hybrid samples populate intermediate regions between the clusters for the three classes. Figure 7(c) shows that the uniform dataset covers a larger space than the simplex dataset.




FIG. 9. ROC-curves for the different GAN-generated datasets from each GAN.

boundaries. The preprocessed detector noise is then split into chunks of length 1,024, the same dimensionality as the GAN input and output. Following this procedure, we accumulate just over 250,000 background samples in total. Since the GAN training data and output are scaled between $[-1, 1]$, the samples are scaled to a signal-to-noise ratio (SNR) ratio that is sampled uniformly on $U[1, 16]$ before injecting them into the preprocessed detector noise. This is done by first computing ρ_{opt} for each generated sample according to

$$\rho_{\text{opt}}^2 = 4 \int_{f_{\min}}^{f_{\max}} \frac{|\hat{h}(f)|^2}{S_n(f)} df, \quad (10)$$

where $\hat{h}(f)$ and $S_n(f)$ are the Fourier transform of the input sample (blip, tomte, BBH) and the detector noise power spectral density (PSD), respectively [66] (which we set to unity for convenience as we are working with whitened data). The sample can then be scaled to the desired ρ_{opt} on $U[1, 16]$.

For each GAN, we train three separate CNNs on the three generated datasets (vertex, simplex, uniform). For the vertex dataset, the three different vertex locations in the class space are sampled with equal probability. For the uniform and simplex datasets, samples are drawn uniformly from their respective spaces.

Each training dataset comprises 100,000 samples, with 50% glitch/signal plus Gaussian noise and 50% Gaussian noise only. We test each CNN on a real data subset comprising 7,500 samples (2,500 from each class), sampled randomly from a distribution of 20,850 samples (6,950 from each class). This results in a test dataset of 15,000 samples in total (50% LIGO noise only) in each iteration of training/testing. We ensure that none of the GAN training data appear in the test set, although they are taken from the same distribution.

The results are presented using the area-under-the-curve (AUC) metric. AUC is a metric commonly used in machine learning to evaluate the performance of a classification model, particularly in the context of binary classification problems. It is associated with the receiver operating characteristic (ROC) curve (see Fig. 9), which is a graphical representation of the trade-off between true positive rate (sensitivity) and false positive rate (1-specificity) across different thresholds. It provides a single metric between

0 and 1 that summarizes the overall discriminatory power of a model across different classification thresholds. We repeat the experiment 5 times and report the mean results, randomly generating the training data and randomly sampling the backgrounds and test data in each iteration. We keep the real test datasets constant for each GAN-generated dataset in each iteration.

4. Fitting-factor study

We investigate the accuracy of cDVGAN's BBH signals under a fitting-factor study, showing that most of the synthetic signals are consistent with the original signals simulated with *IMRPhenomD*. The fitting-factor of a signal is defined as the maximum match of that signal over the templates of a template bank [67].

While cDVGAN's signals are used as the signal injections in a matched-filter search, the template bank is created using *mbank* [68], and the corresponding templates are simulated using *IMRPhenomD*. We aim for good coverage of the parameter space of cDVGAN's training signals ($30 \leq m_2 \leq m_1 \leq 160$) beyond the 97% requirement for matched-filter standards, using a holdout dataset from the cDVGAN BBH distribution to validate the template bank.³ We sample *mbank*'s normalizing flow until the parameter space is covered evenly with an average fitting factor of over 99% (only approximately 5% of signals yield a fitting factor of under 99%). This yields 1,500 templates, allowing for a thorough exploration of the parameter space.

We calculate the fitting factor in the time domain on a fixed time grid since the training signals and generated signals are truncated at 0.25s around the merger time. We compute the fitting-factor over the 2,500 GAN training samples simulated using *IMRPhenomD* and compare it to the fitting factor computed on 2,500 synthetic signals from cDVGAN.

An important caveat to highlight is that the class-conditional approach featured in cDVGAN is focused towards a glitch generator application rather than a waveform generator for GW searches. To achieve better accuracy in the latter case, it would be better to modify the cDVGAN architecture to learn BBH signals only. This would simplify the training schedule to learn one

³We validate the bank using a flat PSD, required for time-domain match calculations using *mbank*.

distribution of signals rather than three distributions in one model. Furthermore, the model could also be conditioned on continuous source parameters rather than discrete class information, giving the user control over the parameter space they wish to generate signals in.

IV. RESULTS

A. Training with GAN data

The AUC values for each CNN model over the entire SNR range (1–16) are shown in Table I. We also scale the test sets between SNRs of 1–8 and 8–16 and record the AUC of the same CNNs to investigate performance for quieter and louder samples, which can be seen in Tables II and III, respectively. Examples of ROC curves from one of the testing iterations are shown in Fig. 9. All tables show that the simplex and uniform datasets from cDVGAN yield the highest AUC results for CNNs, with the simplex dataset yielding the highest overall AUC.

Ablation studies. The ablation between cDVGAN and cWGAN in Table I shows that the adversarial feedback from the first-order derivative discriminator improves the features of the synthetic data in all three datasets. This

suggests that the first-order derivative discriminator was successful in improving the GAN output. cDVGAN2 was unsuccessful in improving upon cDVGAN’s results, yielding a slightly lower performance, but improves upon the performance of cWGAN’s simplex and uniform datasets, particularly for the higher SNR range (Table III).

The AUC performance yielded from McDVGANn’s datasets is competitive among the GAN models. Both its simplex and uniform datasets yield better AUC performance than McGANn’s counterparts, with the simplex dataset yielding the highest overall AUC performance between the two models. The vertex datasets from McGANn and McDVGANn yield comparable performance. This suggests that the derivative discriminator can be effective in a more traditional GAN architecture.

These results show that incorporating derivative discriminators can improve the synthetic data in multiple GAN architectures and indicate that analyzing first-order derivatives in a separate auxiliary discriminator is superior to using both first and second-order derivative discriminators for modeling the dataset covered in this study (in the cWGAN architecture). For all GAN models, hybrid datasets provide the best overall AUC performance for the

TABLE I. The area-under-curve (AUC) yielded on a real test set by CNNs trained on each synthetic dataset from each GAN. The results represent the mean AUC over five iterations, where the bounds are calculated using the standard deviations over the five iterations. The best result overall is shown in bold text.

Dataset	cDVGAN (ours)	cDVGAN2 (ours)	cWGAN	McGANn	McDVGANn (ours)
Vertex trained	0.771 ± 0.012	0.758 ± 0.008	0.762 ± 0.018	0.768 ± 0.016	0.768 ± 0.022
Simplex trained	0.802 ± 0.019	0.789 ± 0.010	0.788 ± 0.009	0.759 ± 0.010	0.786 ± 0.012
Uniform trained	0.797 ± 0.022	0.791 ± 0.009	0.777 ± 0.014	0.770 ± 0.014	0.778 ± 0.012

TABLE II. The AUC for test samples under an SNR of 8. The results represent the mean AUC over five iterations, where the bounds are calculated using the standard deviations over the five iterations. The best result overall is shown in bold text.

Dataset	cDVGAN (ours)	cDVGAN2 (ours)	cWGAN	McGANn	McDVGANn (ours)
Vertex trained	0.689 ± 0.009	0.680 ± 0.006	0.685 ± 0.017	0.687 ± 0.010	0.668 ± 0.020
Simplex trained	0.698 ± 0.013	0.686 ± 0.007	0.695 ± 0.010	0.673 ± 0.008	0.676 ± 0.010
Uniform trained	0.702 ± 0.014	0.693 ± 0.008	0.692 ± 0.009	0.680 ± 0.013	0.671 ± 0.006

TABLE III. The AUC for samples above an SNR of 8. The results represent the mean AUC over five iterations, where the bounds are calculated using the standard deviations over the five iterations. The best result overall is shown in bold text.

Dataset	cDVGAN (ours)	cDVGAN2 (ours)	cWGAN	McGANn	McDVGANn (ours)
Vertex trained	0.844 ± 0.015	0.827 ± 0.010	0.830 ± 0.022	0.841 ± 0.021	0.856 ± 0.027
Simplex trained	0.893 ± 0.029	0.881 ± 0.013	0.867 ± 0.010	0.836 ± 0.012	0.885 ± 0.013
Uniform trained	0.883 ± 0.031	0.878 ± 0.011	0.851 ± 0.020	0.852 ± 0.020	0.873 ± 0.014

TABLE IV. AUC values over three SNR ranges for different proportions of real:synthetic samples for a training set fixed at 100,000 samples. The results are represented by the mean AUC and standard deviation over five iterations.

SNR	100:0	75:25	50:50	25:75	0:100
1–16	0.900 ± 0.001	0.898 ± 0.002	0.893 ± 0.002	0.887 ± 0.002	0.802 ± 0.019
1–8	0.799 ± 0.001	0.792 ± 0.005	0.787 ± 0.003	0.777 ± 0.003	0.698 ± 0.013
8–16	0.988 ± 0.001	0.987 ± 0.001	0.987 ± 0.001	0.985 ± 0.001	0.893 ± 0.029

CNN. This suggests that GAN-generated hybrid samples are useful for searching for multiple classes of real data when obscured by the detector background. This might offer interesting applications in glitch searches, particularly for those with no clear correlation to auxiliary channels. Such glitches could be conditioned into cDVGAN to generate hybrid samples specific to a subset of LIGO glitch classes for use in a glitch detection algorithm. Combining all three GAN-generated datasets for training CNNs may improve upon these results yet again, although this is left to future work.

B. Combining real and cDVGAN data for improved training

In this section, we augment real datasets with GAN-generated data to improve upon the classification performance in the previous section. We augment the real data with the simplex dataset from cDVGAN since it yields the best performance. We compare cDVGAN data with traditional duplication of real training samples for data augmentation.

Maintaining a real hold-out test set of 7,500 samples (2,500 from each class) as in Sec. IV A, we use all remaining samples from the real distribution for training, which amounts to 13,350 samples (4,450 from each class). Fixing the training data size again at 100,000 (50,000 glitch/signal + noise samples, 50,000 noise-only samples), we vary the proportion of real and GAN-generated samples. Since there is no limit to generating cDVGAN data, we duplicate the real training data before injection to reach the required number of samples. This is done to control the effects of the background noise on the training of CNNs, since CNNs are sensitive to the number of different backgrounds seen during training.

The results in Table IV show that the performance drops only slightly with smaller proportions of real data. The CNN performance remains competitive even with only 25% of real training samples with only a 1% drop in overall AUC performance compared to only using real data. The second and third rows of Table IV show that this decrease in AUC performance occurs mostly for lower SNR (<8) samples. The decrease in performance for louder samples is minimal (<1% decrease). This indicates that the synthetic data are competitive for augmenting the training data for CNNs when including a relatively small amount of real data in the training set. Although there is a more substantial decrease in performance when using 100% synthetic

samples, this might be improved by including other synthetic datasets from cDVGAN in the training schedule.

C. Fitting-factor results

The results indicate that cDVGAN’s BBH signals are generally consistent with *IMRPhenomD*. The search calculates an average fitting-factor of 0.994 ± 0.0423 for the 2,500 GAN training signals with fifth and tenth percentiles at 0.972 and 0.994, respectively. Conversely, the experiment yields an average fitting-factor of 0.976 ± 0.045 for 2,500 synthetic cDVGAN signals with fifth and tenth percentiles at 0.893 and 0.951, respectively.

Figure 10 shows a histogram of the fitting-factors yielded from the real and synthetic signals and the distribution of the recovered parameters of both datasets. Figure 10(a) shows the fitting-factor distribution of cDVGAN’s signals is similar to that of the real training signals, although a minor decrease in accuracy is observed. Figure 10(b) shows that cDVGAN’s signals cover most of the parameter space well, although there is a slight mismatch in the lower ends of both the $M1$ and $M2$ distributions.

Although the accuracy of cDVGAN’s signals is lower than that of *IMRPhenomD*, the quality of cDVGAN’s signals and coverage of the parameter space could be improved with further training, or by considering a different modeling approach as outlined in Sec. III C 4. Furthermore, cDVGAN is competitive with state-of-the-art surrogate models in terms of inference speed on a CPU [69,70], and far surpasses them with the use of a GPU. For example, cDVGAN can generate 2,500 signals of the learned BBH class in approximately 18 s on a CPU, and only 0.04 s on a GPU.




FIG. 10. Plots of fitting-factors and corresponding best-fit template parameters for real GAN training signals and synthetic cDVGAN signals.

V. CONCLUSION

Time-domain generative modeling in the GAN framework has shown potential to improve GW data analysis. Using GANs to learn distributions of GWs and other events of interest such as detector glitches can be useful for data augmentation tasks, validating detection schemes for unmodeled waveforms such as [32,71,72], or be used to construct mock data challenges. This work presents a novel conditional GAN, called cDVGAN, for generating distinct time-domain classes, including two classes of unmodeled glitches and one class of modeled BBH signals. cDVGAN uses additional adversarial feedback on the first-order derivatives of training samples in an auxiliary discriminator and generates realistic samples that span the variation within each class. It also allows for the explicit control of the mixing of classes. Thus, it is capable of generating generalized hybrid samples that are outside of the limited training distribution and span the variation between classes by sampling the continuous class space.

We use ablation studies to show the effectiveness of using auxiliary discriminators to analyze sample derivatives in an experiment that uses GAN-generated data to train convolutional neural networks (CNNs) to detect real samples in LIGO detector noise. An ablation study between cDVGAN and its vanilla cWGAN counterpart shows that the additional adversarial feedback from the first-order derivative discriminator yields generated data that is more useful for training a CNN detection algorithm.

The ablation study between cDVGAN2 and cDVGAN reveals that second-order derivative discrimination did not improve the performance under this problem scheme, although the performance of cDVGAN2 is competitive with other baseline GANs. Another ablation study between our McDVGANn model and McGANn [51] indicates that the method can be effective under a traditional cGAN architecture. These results suggest that providing adversarial feedback on derivatives on top of the original samples can improve the learning of GANs on continuous time series, and in particular, events of interest to the GW physics community.

Furthermore, our experiments demonstrate the effectiveness of GAN-generated hybrid samples for training detection algorithms. The best overall synthetic dataset for training CNNs was cDVGAN's simplex dataset, while hybrid datasets from other GANs yielded better training sets than the standard vertex dataset. We also combine GAN-generated data with real data to improve the performance of CNN models for glitch and signal searches, showing cDVGAN as a viable approach for data augmentation. Lastly, we implement a fitting-factor study that shows cDVGAN's BBH signals are consistent with the *IMRPhenomD* waveform routine used to generate the cDVGAN training signals. Although there are some inconsistencies and a small decrease in accuracy, the synthetic signals generally match well with a template bank

for the corresponding parameter range. The cDVGAN signals have good coverage on most of the parameter range and can be generated very efficiently, particularly with the use of a GPU.

Since hyperparameter optimization was not the focus of this research, investigations could be made into better architectures for cDVGAN. For example, optimization of the η [Eq. (8)] hyperparameters controlling the contribution of each discriminator to the generator loss might yield better generated data. Including a consistency term, as in [73], may also improve the generated data from cDVGAN. Expanding cDVGAN to other representations of the data, might also improve the quality of the generated data. Finally, research into better CNNs or other detection algorithms that can make use of the GAN-generated data might also result in efficient and scalable analysis solutions towards the next-generation detectors.

Extending cDVGAN to other glitch types is vital to significantly stimulate GW data analysis. This study takes a step towards this goal, showing how arbitrary time-domain glitches or signals can be conditioned into one generative model. Constructing time-domain representations of unmodeled glitches is challenging, but made possible using algorithms such as *BayesWave* to isolate them from the detector background. Once accurate glitch representations of other LIGO glitch types are constructed, cDVGAN is scheduled for further development in a next-generation glitch generator. Covering the entire LIGO glitch space with cDVGAN will result in a model more representative of LIGO glitches and a useful tool for downstream analysis.

ACKNOWLEDGMENTS

This research was conducted within the ET Technologies project (Project No. PROJ-03612), which is partly funded by EFRO, the Province of Limburg and The Dutch Ministry of Economic Affairs and Climate Policy within the REACT-EU Programme of OP Zuid. The authors are grateful for contributions by members of the ET Technologies research team, in particular, Stefano Schmidt, Andrew Miller, and Sarah Caudill. 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 the National Science Foundation Grants No. PHY-0757058 and No. PHY-0823459.

APPENDIX

1. cDVGAN architecture

Table V below details cDVGAN's architecture, comprising a discriminator, derivative discriminator and a generator.

TABLE V. The architecture and hyperparameters describing cDVGAN, which consists of a base discriminator, a derivative (DV) discriminator and a generator convolutional network. The additional discriminator of cDVGAN2 follows the same architecture as the DV discriminator but with an input shape of 1022. The number of parameters (param.) in each model are shown in brackets beside the name of each model component.

Discriminator (3.5M param.)					
Operation	Output shape	Kernel size	Stride	Dropout	Activation
Input	(1024)	0	...
Reshape	(64, 16)	0	...
Convolutional	(64, 128)	14	2	0.5	Leaky ReLU
Convolutional	(32, 128)	14	2	0.5	Leaky ReLU
Convolutional	(16, 256)	14	2	0.5	Leaky ReLU
Convolutional	(8, 256)	14	2	0.5	Leaky ReLU
Convolutional	(4, 512)	14	2	0.5	Leaky ReLU
Global avg. pooling	(512)	0.5	...
Avg. pooling dense	(128)	0.2	Leaky ReLU
Dense	(1)	0	Linear
Class input	(3)
Class dense	(128)	0	Linear
Scalar product	(1)
Dense + scalar product	(1)
DV Discriminator (1.1M param.)					
Operation	Output shape	Kernel size	Stride	Dropout	Activation
Input	(1023)	0	...
Dense	(512)	0	Leaky ReLU
Reshape	(32, 16)	0	...
Convolutional	(32, 64)	5	2	0.5	Leaky ReLU
Convolutional	(16, 128)	5	2	0.5	Leaky ReLU
Convolutional	(8, 256)	5	2	0.5	Leaky ReLU
Convolutional	(4, 256)	5	2	0.5	Leaky ReLU
Global avg. pooling	(256)	0.5	...
Avg. pooling dense	(128)	0.2	Leaky ReLU
Dense	(1)	0	Linear
Class input	(3)
Class dense	(128)	0	Linear
Scalar product	(1)
Dense + scalar product	(1)
Generator 3.5M param.					
Operation	Output shape	Kernel size	Stride	BN	Activation
Latent input	(100)	\times	...
Class Input	(3)	\times	...
Class Dense	(32)	\times	...
Concatenate	(132)	\times	-
Dense	(1024)	\times	ReLU
Reshape	(32, 32)	\times	...
Transposed conv.	(64, 512)	18	2	\checkmark	ReLU
Transposed conv.	(128, 256)	18	2	\checkmark	ReLU
Transposed conv.	(256, 128)	18	2	\checkmark	ReLU
Transposed conv.	(512, 64)	18	2	\checkmark	ReLU
Transposed conv.	(1024, 1)	18	2	\times	Linear
Flatten	(1024)	\times	...
Optimizer	RMSprop ($\alpha = 0.0001$)				
Batch size	512				
Epochs	500				
Loss	Wasserstein				

TABLE VI. The architecture of the CNN (1.2M parameters) used during experiments.

Operation	Output shape	Kernel size	Stride	Dropout	Activation
Input	(1024)	0	...
Reshape	(1024, 1)	0	...
Convolutional	(512, 256)	5	2	0.5	Leaky ReLU
Convolutional	(256, 128)	5	2	0.5	Leaky ReLU
Convolutional	(128, 64)	5	2	0.5	Leaky ReLU
Convolutional	(64, 32)	5	2	0.5	Leaky ReLU
Flatten	(2048)	0.5	...
Dense	(512)	0	Leaky ReLU
Dense	(1)	0	Sigmoid
Optimizer	Adam ($\alpha = 0.001$)				
Batch size	64				
Epochs	20				
Loss	Binary cross entropy				

2. CNN architecture

Table VI details the architecture of the CNN used during experimentation.

3. Vertex, simplex, and uniform datasets

Figures 11–13 show vertex, simplex and uniform samples generated using different sampling methods in cDVGAN's conditioned class vector.

4. Class interpolation

Figure 14 shows how cDVGAN's class vector can be sampled using a k=2 simplex to interpolate between the three learned classes.

5. Gravity spy analysis of GAN training data

As we have discussed in Sec. III C 3, we have scaled some examples of blip and tomte glitches $g(t)$, adding them to the noise $n(t)$ from the third observing (O3) run in the GPS time range [1262540000, 1262540040]. Note that the glitches have been extracted from their original noise following the procedure previously presented in Sec. III B. Afterwards, the time series $s(t) = g(t) + n(t)$ was classified with Gravity Spy, providing a class label and a class confidence c_{GS} .

In Fig. 15, we present the results of classifying three different denoized blips reinjected in O3 noise. In the top panel, we plot the confidence c_{GS} as a function of the optimal SNR ρ_{opt} , defined in Eq. (10). In the middle panel, we show the time-series of the blip as outputted by the

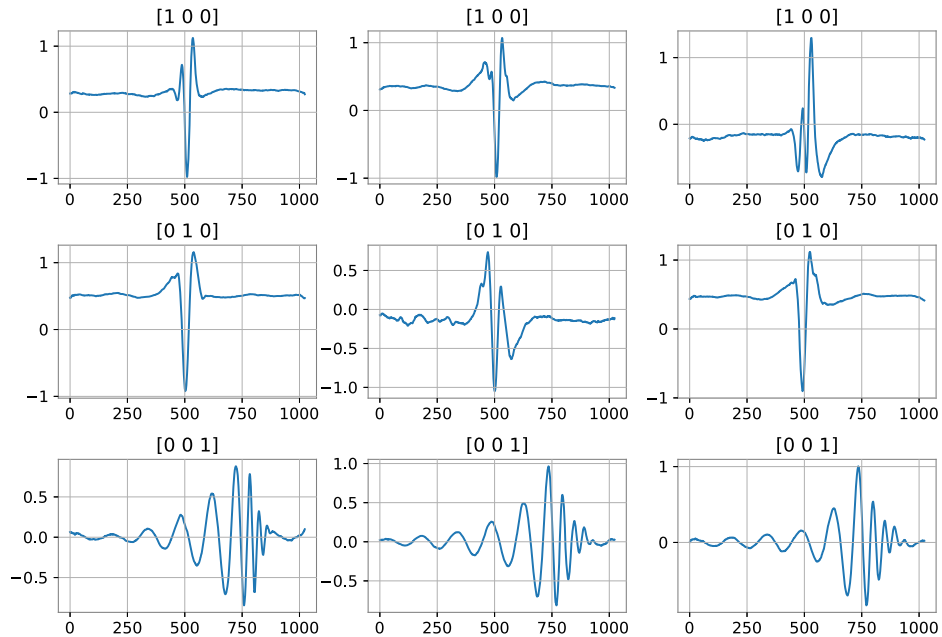


FIG. 11. Standard 3-class vertex generations from cDVGAN.

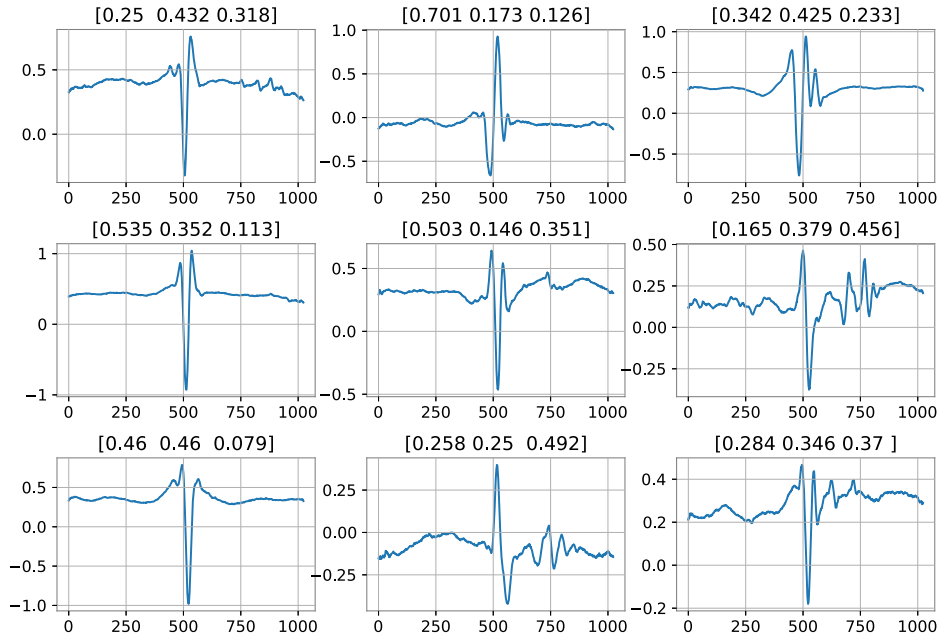


FIG. 12. Simplex generations from cDVGAN.

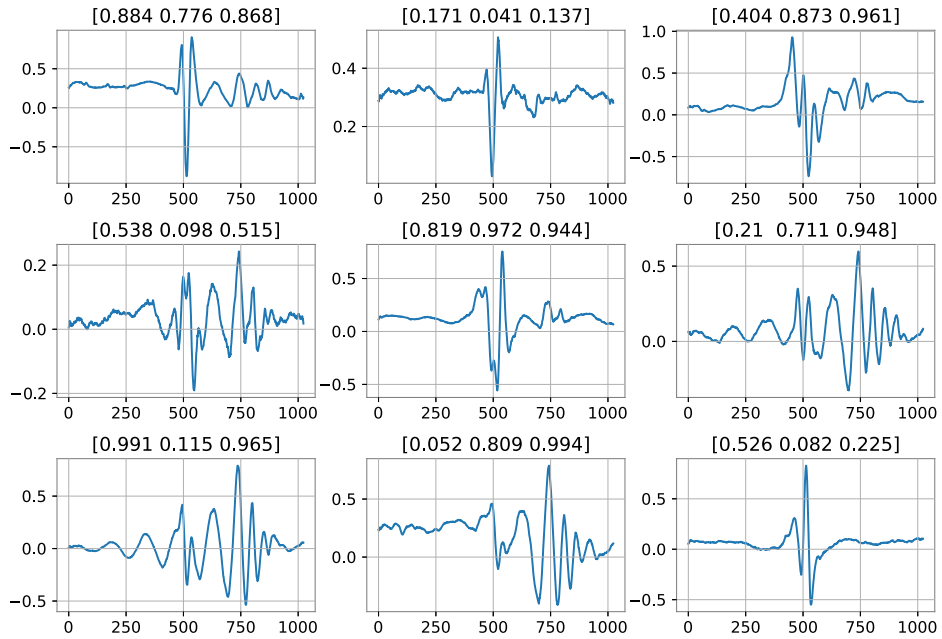


FIG. 13. Uniform generations from cDVGAN.

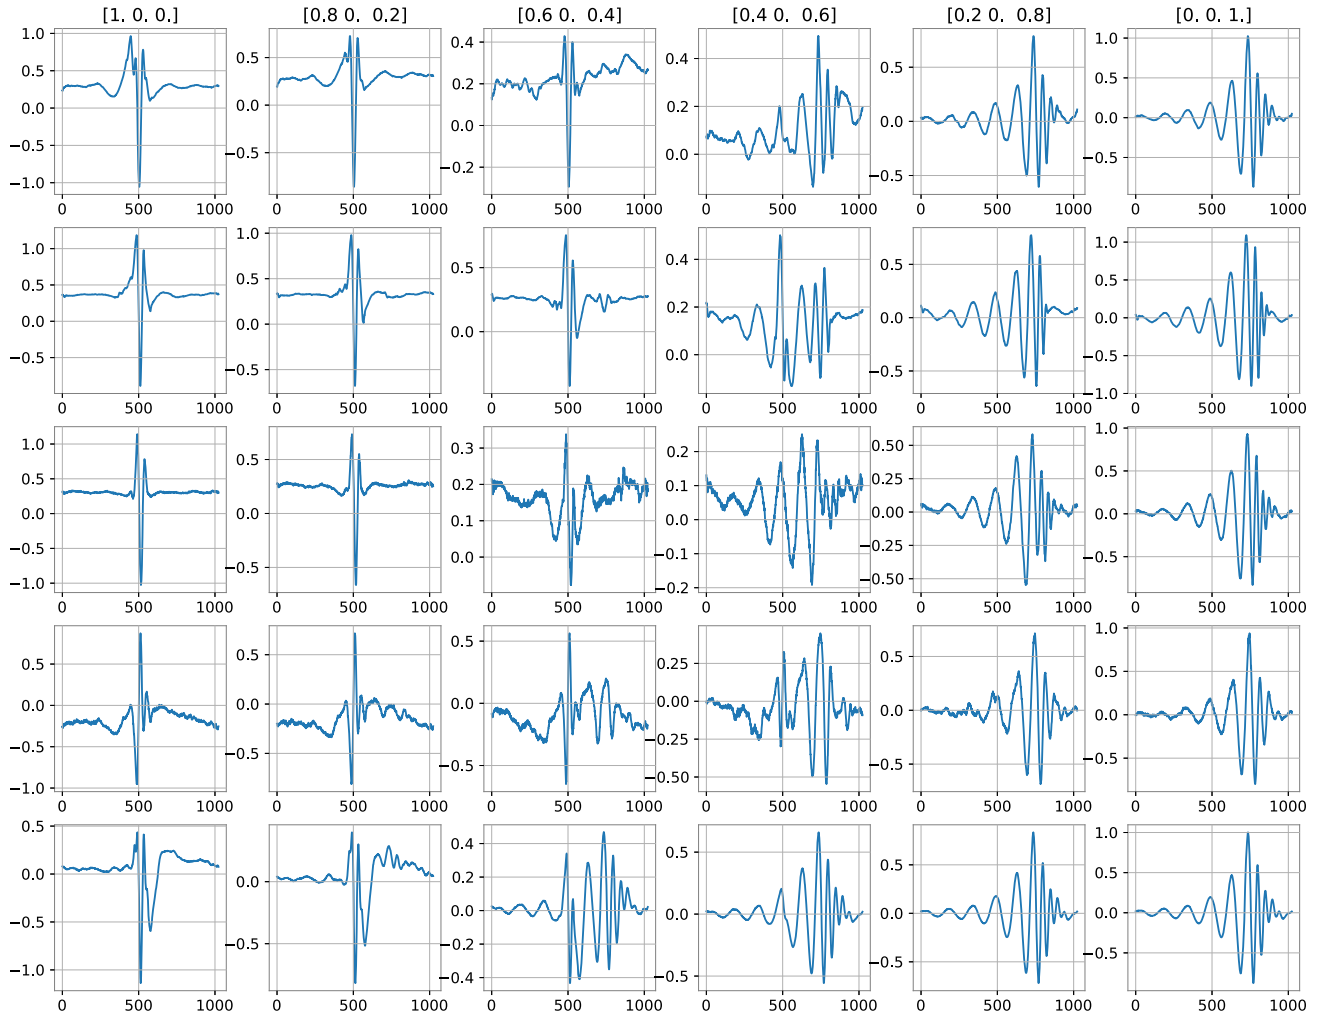


FIG. 14. Interpolation between blip and BBH classes for cDVGAN (first row), cDVGAN2 (second row), cWGAN (third row), McGANn (fourth row), and McDVGANn (fifth row). The class input is shown at the top of each column, while the latent input of the generator is kept constant.

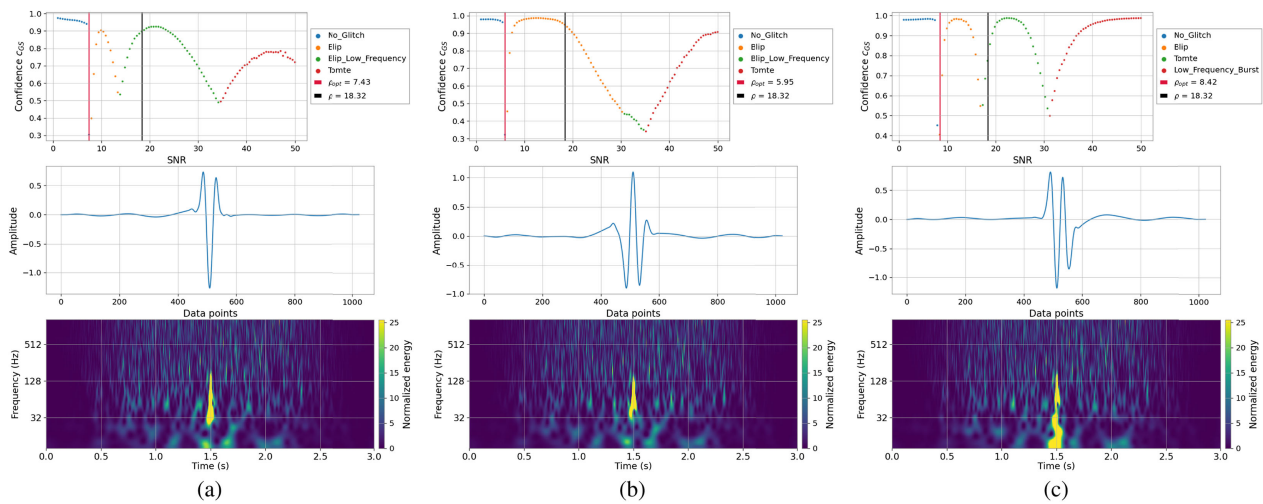


FIG. 15. Top panel: Three examples of blip glitches classified by *Gravity Spy*, with the classification confidences as a function of SNR. Middle panel: The time-series representation of the classified glitch. Bottom panel: The corresponding spectrogram representation of the classified glitch after injecting into whitened detector noise.

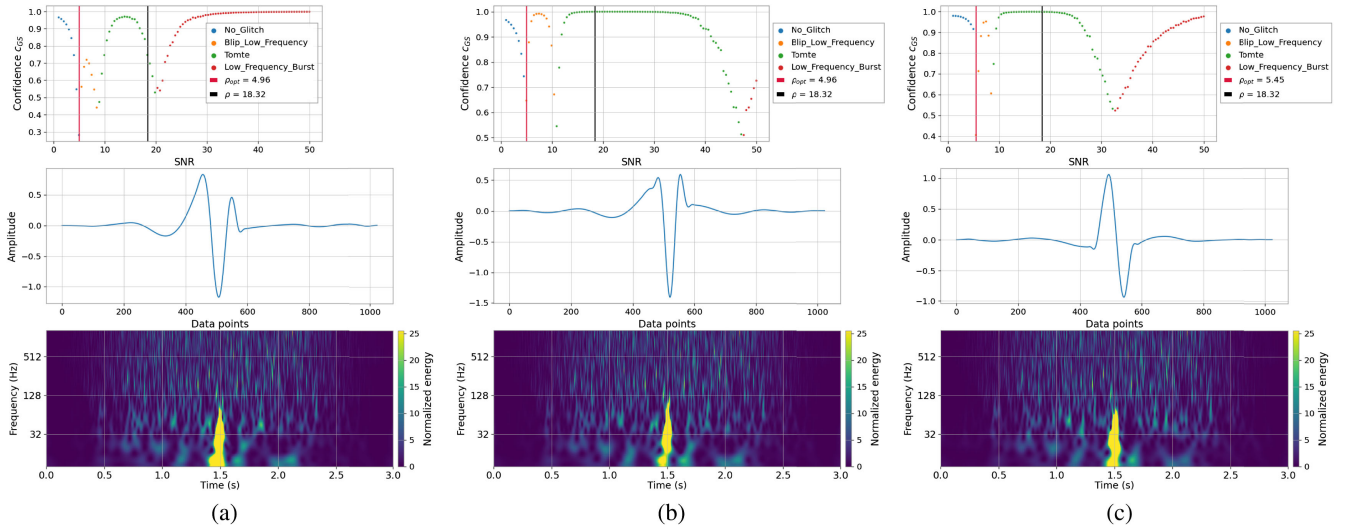


FIG. 16. Top panel: Three examples of tomte glitches classified by *Gravity Spy*, with the classification confidences as a function of SNR. Middle panel: The time-series representation of the classified glitch, Bottom panel: The corresponding spectrogram representation of the classified glitch after injecting into whitened detector noise.

preprocessing, and in the bottom panel, we show the spectrogram or time-frequency representation of the glitch embedded in O3 real noise at $\rho_{\text{opt}} = 18.32$.

We can observe in Fig. 15 that *Gravity Spy* classifies the time series as No_Glitch for $\text{SNR} \lesssim 7$, which is expected as this algorithm only learns glitches with an $\text{SNR} \leq 7.5$. While these examples are labeled as blips at $\text{SNR} \sim 10$, as the SNR increases they get misclassified as blip_low_frequency or low_frequency_burst, which can be explained by the use of bandpass filtering that attenuates the high-frequency contribution of blips.

Similarly, in Fig. 16, we can also observe that the time series gets classified as No_Glitch, but this time at an $\text{SNR} \lesssim 5$. Then, for $5 \lesssim \text{SNR} < 10$, it gets briefly classified as blip_low_frequency, to then be classified as tomte for most of the SNR range, meaning that the features of tomtes are better preserved than the blip class after data preprocessing.

6. McDVGANn additional analysis

Figure 17 shows that the optimum update schedule for cDVGAN is to update the discriminator one time rather than three times for every generator update.




FIG. 17. A comparison of McDVGANn CNN performance trained with a discriminator update schedule of 1 for each generator update and a discriminator update schedule of 3 times for each generator update, which is used for the original McGANn model. These results suggest that McDVGANn overfits the data with the original McGANn update schedule and that using a discriminator update schedule of 1 is superior.

- [1] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] J. Aasi *et al.* (LIGO Scientific Collaboration), Advanced LIGO, *Classical Quantum Gravity* **32**, 074001 (2015).

- [3] F. Acernese *et al.* (Virgo Collaboration), Advanced Virgo: A second-generation interferometric gravitational wave detector, *Classical Quantum Gravity* **32**, 024001 (2015).
- [4] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and Virgo

- during the first and second observing runs, *Phys. Rev. X* **9**, 031040 (2019).
- [5] R. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), GWTC-2: Compact binary coalescences observed by LIGO and Virgo during the first half of the third observing run, *Phys. Rev. X* **11**, 021053 (2021).
- [6] R. Abbott *et al.* (LIGO Scientific, Virgo, and KAGRA Collaborations), GWTC-3: Compact binary coalescences observed by LIGO and Virgo during the second part of the third observing run, *Phys. Rev. X* **13**, 041039 (2023).
- [7] T. Akutsu *et al.*, Overview of KAGRA: Detector design and construction history, *Prog. Theor. Exp. Phys.* **2021**, 05A101 (2020).
- [8] B. P. Abbott *et al.*, Prospects for observing and localizing gravitational-wave transients with Advanced LIGO, Advanced Virgo and KAGRA, *Living Rev. Relativity* **23**, 3 (2020).
- [9] S. Hild, M. Abernathy, F. Acernese, P. Amaro-Seoane, N. Andersson, K. Arun, F. Barone, B. Barr, M. Barsuglia, M. Beker *et al.*, Sensitivity studies for third-generation gravitational wave observatories, *Classical Quantum Gravity* **28**, 094013 (2011).
- [10] E. D. Hall, Cosmic explorer: A next-generation ground-based gravitational-wave observatory, *Galaxies* **10**, 90 (2022).
- [11] F. Iacovelli, M. Mancarella, S. Foffa, and M. Maggiore, Forecasting the detection capabilities of third-generation gravitational-wave detectors using GFAST, *Astrophys. J.* **941**, 208 (2022).
- [12] V. Kalogera *et al.*, *The Next Generation Global Gravitational Wave Observatory: The Science Book* (Gravitational Wave International Committee, 2021).
- [13] C. Fryer and K. New, Gravitational waves from gravitational collapse, *Living Rev. Relativity* **14**, 1 (2011).
- [14] L. Baiotti, I. Hawke, L. Rezzolla, and E. Schnetter, Details on the gravitational-wave emission from rotating gravitational collapse in 3d (2007).
- [15] D. Davis *et al.*, LIGO detector characterization in the second and third observing runs, *Classical Quantum Gravity* **38**, 135014 (2021).
- [16] F. Acernese *et al.* (Virgo Collaboration), Virgo detector characterization and data quality: Results from the O3 run, *Classical Quantum Gravity* **40**, 185006 (2023).
- [17] T. Akutsu *et al.*, Overview of KAGRA: Calibration, detector characterization, physical environmental monitors, and the geophysics interferometer, *Prog. Theor. Exp. Phys.* **2021**, 05A102 (2021).
- [18] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO's first observing run, *Classical Quantum Gravity* **35**, 065010 (2018).
- [19] R. Abbott *et al.* (KAGRA, LIGO Scientific, and Virgo Collaborations), All-sky search for continuous gravitational waves from isolated neutron stars using Advanced LIGO and Advanced Virgo O3 data, *Phys. Rev. D* **106**, 102008 (2022).
- [20] B. Steltner, M. A. Papa, H.-B. Eggenstein, R. Prix, M. Bensch, B. Allen, and B. Machenschalk, Deep Einstein@Home all-sky search for continuous gravitational waves in LIGO O3 public data, *Astrophys. J.* **952**, 55 (2023).
- [21] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), Upper limits on the isotropic gravitational-wave background from Advanced LIGO and Advanced Virgo's third observing run, *Phys. Rev. D* **104**, 022004 (2021).
- [22] B. Steltner, M. A. Papa, and H. B. Eggenstein, Identification and removal of non-Gaussian noise transients for gravitational-wave searches, *Phys. Rev. D* **105**, 022005 (2022).
- [23] C. Pankow *et al.*, Mitigation of the instrumental noise transient in gravitational-wave data surrounding GW170817, *Phys. Rev. D* **98**, 084016 (2018).
- [24] D. Davis, T. Massinger, A. Lundgren, J. C. Driggers, A. L. Urban, and L. Nuttall, Improving the sensitivity of Advanced LIGO using noise subtraction, *Classical Quantum Gravity* **36**, 055011 (2019).
- [25] J. C. Driggers *et al.*, Improving astrophysical parameter estimation via offline noise subtraction for Advanced LIGO, *Phys. Rev. D* **99**, 042001 (2019).
- [26] L. Blackburn *et al.*, The LSC Glitch Group: Monitoring noise transients during the fifth LIGO science run, *Classical Quantum Gravity* **25**, 184004 (2008).
- [27] B. P. Abbott *et al.* (LIGO Scientific and Virgo Collaborations), Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914, *Classical Quantum Gravity* **33**, 134001 (2016).
- [28] B. P. Abbott *et al.*, A guide to LIGO–Virgo detector noise and extraction of transient gravitational-wave signals, *Classical Quantum Gravity* **37**, 055002 (2020).
- [29] B. J. Owen and B. S. Sathyaprakash, Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement, *Phys. Rev. D* **60**, 022002 (1999).
- [30] S. A. Usman *et al.*, The PyCBC search for gravitational waves from compact binary coalescence, *Classical Quantum Gravity* **33**, 215004 (2016).
- [31] K. Cannon *et al.*, Gstlal: A software framework for gravitational wave discovery, *SoftwareX* **14**, 100680 (2021).
- [32] F. Robinet, N. Arnaud, N. Leroy, A. Lundgren, D. Macleod, and J. McIver, Omicron: A tool to characterize transient noise in gravitational-wave detectors, *SoftwareX* **12**, 100620 (2020).
- [33] H. Gabbard, M. Williams, F. Hayes, and C. Messenger, Matching matched filtering with deep networks for gravitational-wave astronomy, *Phys. Rev. Lett.* **120**, 141103 (2018).
- [34] T. D. Gebhard, N. Kilbertus, I. Harry, and B. Schölkopf, Convolutional neural networks: A magic bullet for gravitational-wave detection?, *Phys. Rev. D* **100**, 063015 (2019).
- [35] P. G. Krastev, Real-time detection of gravitational waves from binary neutron stars using artificial neural networks, *Phys. Lett. B* **803**, 135330 (2020).
- [36] E. Cuoco *et al.*, Enhancing gravitational-wave science with machine learning, *Mach. Learn.* **2**, 011002 (2020).
- [37] M. Zevin *et al.*, Gravity spy: Integrating Advanced LIGO detector characterization, machine learning, and citizen science, *Classical Quantum Gravity* **34**, 064003 (2017).
- [38] Y. Wu, M. Zevin, C. P. L. Berry, K. Crowston, C. Østerlund, Z. Doctor, S. Banagiri, C. B. Jackson, V. Kalogera, and A. K. Katsaggelos, Advancing glitch classification in

- gravity spy: Multi-view fusion with attention-based machine learning for Advanced LIGO's fourth observing run, [arXiv:2401.12913](https://arxiv.org/abs/2401.12913).
- [39] S. Bahaadini, N. Rohani, S. Coughlin, M. Zevin, V. Kalogera, and A. K. Katsaggelos, Deep multi-view models for glitch classification, *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (New Orleans, USA, 2017), pp. 2931–2935, <https://api.semanticscholar.org/CorpusID:13008908>.
- [40] S. Bahaadini, V. Noroozi, N. Rohani, S. Coughlin, M. Zevin, J. Smith, V. Kalogera, and A. Katsaggelos, Machine learning for gravity spy: Glitch classification and dataset, *Inf. Sci.* **444**, 172 (2018).
- [41] D. George, H. Shen, and E. Huerta, Classification and unsupervised clustering of LIGO data with deep transfer learning, *Phys. Rev. D* **97**, 101501 (2018).
- [42] I. J. Goodfellow *et al.*, Generative adversarial networks, *Adv. Neural Inf. Process. Syst.* **3**, 2672 (2014).
- [43] J. Powell, L. Sun, K. Gereb, P. D. Lasky, and M. Dollmann, Generating transient noise artifacts in gravitational-wave detector data with generative adversarial networks, *Classical Quantum Gravity* **40**, 035006 (2023).
- [44] J. Yan, A. P. Leung, and D. C. Y. Hui, On improving the performance of glitch classification for gravitational wave detection by using generative adversarial networks, *Mon. Not. R. Astron. Soc.* **515**, 4606 (2022).
- [45] J. Abadie *et al.*, Search for gravitational waves from compact binary coalescence in LIGO and Virgo data from S5 and VSR1, *Phys. Rev. D* **82**, 102001 (2010).
- [46] J. Abadie *et al.*, All-sky search for gravitational-wave bursts in the first joint LIGO-GEO-Virgo run, *Phys. Rev. D* **81**, 102001 (2010).
- [47] C. Biwer *et al.*, Validating gravitational-wave detections: The Advanced LIGO hardware injection system, *Phys. Rev. D* **95**, 062002 (2017).
- [48] M. Lopez, V. Boudart, K. Buijsman, A. Reza, and S. Caudill, Simulating transient noise bursts in LIGO with generative adversarial networks, *Phys. Rev. D* **106**, 023027 (2022).
- [49] M. Arjovsky, S. Chintala, and L. Bottou, Wasserstein generative adversarial networks, in *Proceedings of the 34th International Conference on Machine Learning*, Proceedings of Machine Learning Research Vol. 70, edited by D. Precup and Y. W. Teh (PMLR, 2017), pp. 214–223.
- [50] N. J. Cornish and T. B. Littenberg, Bayeswave: Bayesian inference for gravitational wave bursts and instrument glitches, *Classical Quantum Gravity* **32**, 135012 (2015).
- [51] J. McGinn, C. Messenger, M. J. Williams, and I. S. Heng, Generalised gravitational wave burst generation with generative adversarial networks, *Classical Quantum Gravity* **38**, 155005 (2021).
- [52] I. Gulrajani *et al.*, Improving the improved training of Wasserstein GANs: A consistency term and its dual effect, in *International Conference on Learning Representations (ICLR)* (Vancouver, CA, 2018), https://openreview.net/forum?id=Hk4_qwRb.
- [53] T. Salimans *et al.*, Improved techniques for training GANs, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc, 2016), pp. 2234–2242.
- [54] T. Karras *et al.*, Progressive growing of GANs for improved quality, stability, and variation, in *International Conference on Learning Representations* (Vancouver, Canada, 2018), <https://openreview.net/forum?id=Hk99zCeAb>.
- [55] M. Mirza and S. Osindero, Conditional generative adversarial nets (2014), <https://api.semanticscholar.org/CorpusID:12803511>.
- [56] T. Miyato and M. Koyama, cGANs with projection discriminator, in *International Conference on Learning Representations* (Vancouver, CA, 2018), <https://openreview.net/forum?id=ByS1VpgRZ>.
- [57] T. Dooney, S. Bromuri, and L. Curier, Dygan: Stabilize Wasserstein GAN training for time-domain gravitational wave physics, in *2022 IEEE International Conference on Big Data (Big Data)* (Osaka, Japan, 2022), pp. 5468–5477, [10.1109/BigData55660.2022.10021080](https://doi.org/10.1109/BigData55660.2022.10021080).
- [58] M. Cabero *et al.*, Blip glitches in Advanced LIGO data, *Classical Quantum Gravity* **36**, 15 (2019).
- [59] J. Glanzer *et al.*, Data quality up to the third observing run of Advanced LIGO: Gravity spy glitch classifications, *Classical Quantum Gravity* **40**, 065004 (2023).
- [60] D. M. Macleod, J. S. Areeda, S. B. Coughlin, T. J. Massinger, and A. L. Urban, GWpy: A Python package for gravitational-wave astrophysics, *SoftwareX* **13**, 100657 (2021).
- [61] R. Abbott *et al.* (KAGRA, Virgo, and LIGO Scientific Collaborations), Open data from the third observing run of LIGO, Virgo, KAGRA, and GEO, *Astrophys. J. Suppl. Ser.* **267**, 29 (2023).
- [62] A. Savitzky and M. J. E. Golay, Smoothing and differentiation of data by simplified least squares procedures, *Anal. Chem.* **36**, 1627 (1964).
- [63] A. Nitz *et al.*, gwastro/pycbc: v2.0.4 release of pycbc (2022).
- [64] LIGO Scientific Collaboration, LIGO Algorithm Library—LALSuite, free software (GPL) (2018).
- [65] H. I. Fawaz *et al.*, Deep learning for time series classification: A review, *Data Min. Knowl. Discovery* **33**, 917 (2019).
- [66] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, and J. D. E. Creighton, Findchirp: An algorithm for detection of gravitational waves from inspiraling compact binaries, *Phys. Rev. D* **85**, 122006 (2012).
- [67] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg, Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries, *Phys. Rev. D* **77**, 104017 (2008).
- [68] S. Schmidt, B. Gadre, and S. Caudill, Gravitational-wave template banks for novel compact binaries, *Phys. Rev. D* **109**, 042005 (2024).
- [69] V. Varma, S. E. Field, M. A. Scheel, J. Blackman, L. E. Kidder, and H. P. Pfeiffer, Surrogate model of hybridized numerical relativity binary black hole waveforms, *Phys. Rev. D* **99**, 064045 (2019).
- [70] M. Tiglio and A. Villanueva, Reduced order and surrogate models for gravitational waves, *Living Rev. Relativity* **25**, 2 (2022).
- [71] M. Drago, S. Klimentenko, C. Lazzaro, E. Milotti, G. Mitselmakher, V. Necula, B. O'Brian, G. A. Prodi,

- F. Salemi, M. Szczepanczyk, S. Tiwari, V. Tiwari, G. V. G. Vedovato, and I. Yakushin, coherent waveburst, a pipeline for unmodeled gravitational-wave data analysis, *SoftwareX* **14**, 100678 (2021).
- [72] S. Klimenko, I. Yakushin, A. Mercer, and G. Mitselmakher, A coherent method for detection of gravitational wave bursts, *Classical Quantum Gravity* **25**, 114029 (2008).
- [73] W. Xiang *et al.*, Improving the improved training of Wasserstein GANs: A consistency term and its dual effect, in *International Conference on Learning Representations (ICLR)* (Vancouver, CA, 2018).